universität wien

MASTERARBEIT

Titel der Masterarbeit

„Epistemic Authority in Context“

A conceptual synthesis of knowledge

Verfasserin

Martha Cassidy-Brinn, B.A.

angestrebter akademischer Grad

Master of Arts (M.A.)

Wien, 2011

Studienkennzahl lt. Studienblatt:	A 066 941
Studienrichtung lt. Studienblatt:	Philosophie
Betreuer:	Univ.-Prof. Dr. Martin Kusch

Eidesstattliche Erklärung

Hiermit erkläre ich,

- dass ich die vorliegende Arbeit selbst verfasst und nur die angegebene Literatur verwendet habe,

- dass ich dieses Masterthema bisher weder im In- noch im Ausland einer Beurteilerin bzw. einem Beurteiler zur Begutachtung in irgendeiner Form als Prüfungsarbeit vorgelegt habe,

- dass diese Arbeit mit der vom Begutachter beurteilten Arbeit übereinstimmt.

Wien, April 2011

Till fire joins wind
it cannot take a step.

Do men know
it's like that
with knowing and doing?

Devara Dasimayya

Contents

Introduction

Words are beneficial tools which have enabled the human species to survive and flourish. Concepts are equally beneficial, for they are expressed through words. People use the concept of danger, for example, when they talk about danger. Talking about danger is beneficial because it meets the need to warn others and to teach them how to avoid threats to survival. The concept of mushroom can be used both to talk about and to find mushrooms, as opposed to other plants or prey.

Concepts, then, serve a *function* in social life. They meet specific needs. Furthermore, the needs they function to meet affect their extension. For example, having a concept of dangerous mushroom meets the need to warn others of inedible as opposed to edible mushrooms and to teach them about same. Yet, one might object that all mushrooms are *potentially* dangerous. If one were to inhale them or to eat several barrels at once, otherwise innocuous mushrooms could prove harmful. We should not, though, allow these remote possibilities to affect the extension of the concept of dangerous mushroom. If we began to refer to all mushrooms as dangerous, we would no longer be able to use the concept in a fruitful way; its capacity to function in useful warnings and lessons would be lost. Given the function of the concept, it is far better to reserve it for mushrooms that are dangerous when eaten the typical human way.

Notice that with this thesis I do not adopt any radical stance about the social construction of truth or objective reality. Human interests do not create the mushroom nor its chemical makeup. Thus, 'dangerous mushroom' refers to the outside world. Nonetheless, the extension of the concept is irreducibly relative to human interests and needs. If different foods made humans sick, the concept of dangerous mushroom would have a different extension. Likewise, if humans could eat anything at all, there would be no more need for the concept of dangerous mushroom at all. Human interests determine how our concepts carve up the world; they do not determine the world itself.[1]

In this thesis I argue that knowledge is another useful concept which meets human needs. It is useful to be able to recognize a certain sort of relationship

[1] This point is not incompatible with the existence of natural kinds. It could be human interests that determine which natural kinds we conceptually recognize and which natural kinds we ignore.

between a person and a proposition, specifically the sort of relationship that should inspire *trust*. Human societies help themselves by drawing attention to that special kind of trustworthiness; they call it 'knowledge'. Having a concept of knowledge helps people navigate others' assertions and their own beliefs, dividing the trustworthy from the untrustworthy. As they make that division, their goal is to trust only the *true*. Their strategy for reaching that goal involves trusting only the *authoritative*. The concept of knowledge then comprises both goal and strategy: a person who knows that *p* is thus a person who *should be trusted*, because her belief that *p* is true and because she has the authority to believe that *p*.

I do not support this thesis by means of conceptual analysis. Conceptual analysis of knowledge begins with an intuitive grasp of the extension of the concept of knowledge and then dissects those intuitions, attempting to extract a list of individually necessary and jointly sufficient conditions that can accurately predict the analyst's intuitive response to any hypothetical example.[2] While the bulk of Anglo-American epistemology has relied on conceptual analysis, its practitioners still await a detailed explanation and defense of their method. To date no such account has met with general acceptance, or even popularity.[3] Neither shall I attempt to remedy this deficit. My thesis suggests, rather, that we may use a new kind of method to go about conceptualizing knowledge: *conceptual synthesis*. The role model for my presentation and application of conceptual synthesis is Edward Craig's work in *Knowledge and the State of Nature: An Essay in Conceptual Synthesis.*[4]

True to its name, a conceptual synthesis must begin by *finding* and *assembling* conditions, rather than analytically *extracting* them from a given concept. After elaborating on the motivation behind his methodological advances (1.a), I clarify Craig's own use of the method in Chapter One, by breaking his process down into four distinct steps. First, he puts forth a hypothesis as to which function the concept of knowledge might serve. Second, he explicates the conditions that any concept serving that function would require. Third, he synthesizes the explicated conditions into one concept and finally, he compares the synthesized concept to our intuitive

[2] Grundmann 2008, pp. 10 – 11.
[3] The dubiousness of basing the method on intuitions has drawn particularly severe critique and no satisfying response. (Williamson, 2007, p. 215) For specific criticisms, see Bishop & Trout 2004; Buckwalter & Stich forthcoming; Weinburg, Nichols, & Stich 2001; Stich 1990. For an attempted response, see Sosa 2005.
[4] Craig 1990.

grasp of the concept of knowledge. Thus conceptual synthesis ends where conceptual analysis begins. Rather than grounding the entire proceeding, intuitions serve as a final test of the hypothesis. If the comparison reveals no unacceptable discrepancy, then the synthesized concept may be identified with the concept of knowledge. These four steps, hypothesis, explication, synthesis and comparison, comprise the *practical explication* (1.b). After explaining this aspect of the method, I present the results which Craig achieves by using it. Craig hypothesizes that the concept of knowledge might meet the need for inquirers to label good informants (1.c) and accordingly synthesizes a concept of the good informant (1.d).

As Craig acknowledges, his original hypothesis does not prove successful; the final comparison reveals significant discrepancies between the concept of the good informant and the concept of knowledge. To remedy the mismatch, he introduces a strategy which has provoked widely varying responses. Following Bernard Williams, I refer to this aspect of the method as a *genealogy* and devote Chapter Two to its clarification. A genealogy is a story that explains a process of change. First I explain the form it takes in Craig's work; Craig tells a story about the process of objectivization (2.a). His good informant hypothesis describes the needs of a fairly isolated individual, yet all real individuals are embedded in their society. With the story of objectivization, Craig describes how concepts gradually change as people adapt to social demands. This story should bridge the gap between the overly subjective concept of the good informant and the objective concept of knowledge, thus saving Craig's hypothesis. I go on to present the various forms in which his genealogy has been received (2.b), delineating four kinds of genealogies in and inspired by Craig's and Williams's work: rhetorical, imaginary, conjectural, and factual. I also present a fifth approach, exemplified by Klemens Kappel, who omits the genealogy altogether and concentrates on the practical explication alone (2.c). This fifth approach need not entail a total rejection of objectivization; it rather recasts objectivization as a modification of the original hypothesis rather than the synthesized concept. That is, the original hypothesis could be reformulated to take social circumstances into account from the first step on. My own work falls under the fifth approach.

While Chapters One and Two mostly restrict themselves to the presentation of Craig's conceptual synthesis and its reception, the remainder of the thesis consists of

my own application of the method. I suggest a new hypothesis (3.a): the concept of knowledge may function to meet the need for epistemic authority. The alternative hypothesis modifies Craig's in that it *begins* by recognizing the social demands on any public concept, obviating the need for a genealogical story of objectivization. Yet the alternative hypothesis should also include Craig's original good informant hypothesis, rather than replace it (3.b). These two features of the hypothesis forebode trouble (3.c): how may the social and individual demands coexist without cancelling each other out?

To explicate conditions from the hypothesis, I draw inspiration from the work of John Hawthorne and Jason Stanley, hypothesizing that the need for epistemic authority may be derived from the demands of rational practical reasoning (4.a). Subjects need to identify when it is appropriate to believe a proposition, so that they may rationally use that proposition as a premise in their practical reasoning. The concept of epistemic authority could function to label those subjects upon whose authority it is appropriate to believe. From this hypothesis, I explicate the condition that the concept of knowledge must be sensitive to features that affect an inquirer's practical environment, features such as stakes and the end of inquiry (4.b). From this condition I draw the conclusion that, in order to serve my hypothesized function, the concept's extension must be sensitive to the practical environment of the attributor (4.c). Thus, despite the similarities to Hawthorne and Stanley's subject-sensitive invariantism, my synthesis must allow for sensitivity to an attributor's context. The demand for contextualism I practically explicate, however, is for *functional* contextualism, which is neutral towards the question of whether contextual shift is semantic or pragmatic. Thus, I bypass the traditional debate over contextualism and concentrate rather on developing a clearer picture of the behavior of the concept of knowledge, regardless of whether that behavior is due to semantic or pragmatic features. Finally, I return to the problem anticipated in Chapter Three, showing two conflicts between the public demands of the hypothesis and the private demands of the explicated contextualist condition (3.g). Altogether, I explicate three conditions which any concept that met the need for epistemic authority would have to satisfy (3.h), one to meet the need for contextualist sensitivity and two to meet the need for public stability.

Using those three desiderata, I evaluate various conceptions of knowledge in Chapter Four to judge whether they provide a synthesis that would be able to serve my hypothesized function. First I consider two forms of invariantism, high standards (5.a) and moderate (5.b), and two forms of contextualism, salience-sensitive (5.c) and pragmatic (5.d). None of these conceptions manages to synthesize all three conditions into one coherent concept. I then turn to a more successful approach, convention contextualism (5.e), which does satisfy all three desiderata and so does provide a synthesis of the concept of knowledge that would be able to serve my hypothesized function. As my final synthesis of the concept of knowledge (5.f) describes how *knowledge* functions as a label of *authority* in *practical* reasoning, sensitive to *conventional contexts*, I abbreviate the final synthesis KAP–CC.

In Chapter Six, I compare the KAP–CC synthesis to generally shared intuitions about the concept of knowledge, considering several possible discrepancies between how they describe disagreement (6.a), behavior evaluation (6.b), belief (6.c), and useless information (6.d). Finally, I suggest several ways in which further research could widen and deepen my results through normative stipulation and empirical data, including factual and conjectural genealogies (6.e).

1. The Concept of the Good Informant

In the following two chapters, I shall introduce Craig's methodological innovations. As described in the introduction, conceptual analysis proceeds by analyzing the intuitive extension of the concept. Craig begins instead by asking, *what need might having that concept meet?* With a hypothesized need before him, he goes on to ask, *which conditions would a concept that met that need require*? Thus, he does go on to offer an analysis, but the object of the analysis is a *need* rather than an intuitive concept. He analyzes the need to explicate the conditions of its satisfaction. He then puts those conditions together into a synthesis and compares the result to our intuitive concept. If the comparison shows a striking similarity, the hypothesis is likely to be true.

We can best understand Craig's conceptual synthesis as divisible into two distinct parts that I shall present in turn: a practical explication and a genealogy. This chapter is devoted to the practical explication and the results Craig draws from its application.

a. The function of the concept of knowledge

i. Why do people attribute knowledge?

Before presenting the details of his method, we should consider *why* Craig developed it and why he found it appropriate in particular for understanding the concept of knowledge. Craig's motivation is based in an observation and a question.

> There seems to be no known language in which sentences using 'know' do not find a comfortable and colloquial equivalent. The implication is that it answers to some very general needs of human life and thought, and it would surely be interesting to know which and how.[5]

Craig observes that every linguistic community has some close translation for the word "knowledge" and its cognates. This observation inspires the question, why does everyone talk about knowledge? What need do all these people have in common that the concept of knowledge functions to meet? Thus, the widespread use of the word knowledge motivates the first step of the conceptual synthesis, the proposal of a hypothesis as to what function the concept of knowledge might serve.

[5] Craig 1990, p. 2.

It is of course to be taken for granted that it will still be possible to use the concept of knowledge in ways secondary to the hypothesized function. The imperative "Stop!" primarily meets the need of getting someone to stop doing something, but it can also be used to scare someone, to make a joke, or in make-believe. I will assume that these needs are secondary to the first, for if the concept didn't usually function to serve the primary need, then the secondary needs couldn't be met either. If we find such a primary function for knowledge, then we should be able to plausibly characterize its other functions as secondary to that.[6]

One might argue that the primary function of the concept of knowledge is simply to talk about people who have knowledge. Knowledge is out there; the concept of knowledge functions simply to describe it.[7] Craig rejects this answer, reasoning that knowledge is not a natural kind.[8] That is, the phenomenon of knowledge is grounded in the human need for the concept and not the other way around. Yet even if knowledge were a natural kind, this answer would not preclude the applicability of conceptual synthesis to knowledge. After all, whether phenomena of knowledge naturally exist or not, the definition of knowledge is *controversial*, and that controversy is rooted in disagreement over the very standards that its definition must meet. We may approach that controversy by asking, *why do we need the concept of knowledge?* Perhaps the question could be rephrased as, *why do we need to talk about this natural kind?* Nonetheless, there is no immediate reason to assume that the answers will be any less important to quelling the controversy that surrounds the standards for a definition of knowledge.[9]

ii. Value

If the concept of knowledge serves a widespread social function, it has social value. Nonetheless, it is not immediately obvious that Craig's work may be classified as a

[6] (Kappel, 2010, pp. 71 - 72)
[7] Hilary Kornblith offers this response to Craig in Kornblith forthcoming.
[8] Craig 1990, p. 3.
[9] For example, Kornblith assumes that the concept of knowledge need not be immune to outlandish counterexamples. (Kornblith 2002, p. 69.) With a conceptual synthesis, he could defend that assumption, arguing that we do not *need* the concept to be so immune. As is stands, the assumption is not adequately founded and his claim that biology should determine the relevant standards seems capricious.

study of the value of knowledge.[10] Craig's method concerns the function of applications of the concept of knowledge, not knowledge itself.

> We are asking not so much, when is the ascription of a certain concept correct, but rather, why it is applied?[11]

Applications of the concept of knowledge generally take the form of knowledge *attributions*: one subject attributes knowledge to another subject.

By bringing the function of attributions to the fore, Craig analyzes the value of having a concept of a certain phenomenon, rather than the value of the phenomenon itself. The two kinds of values are not identical. For example, having a concept of badness might be useful in so far as it functions to help humans avoid bad things. The concept of badness is valuable, but we need not conclude that bad things are valuable. Thus, the answer to Craig's question need not necessarily reveal anything about the value of knowledge per se.

It would of course defy most epistemologists' intuitions if knowledge per se turned out to be noxious or neutral. In fact, Craig's work does not lead us to such a conclusion. He does argue that the value of the concept of knowledge is based upon the recognition of something valuable. Nonetheless, it is important to keep in mind in what follows that we are concentrating on the function played by knowledge attributions, rather than knowledge per se.

b. Practical explication

A conceptual synthesis begins with a hypothesized need and goes on to synthesize a concept which would meet that need. As Craig puts it, we should take,

> some prima facie plausible hypothesis about what the concept of knowledge does for us, what its role in our life might be, and then ask what a concept having that role would be like, what conditions would govern its application.[12]

If the concept that would meet our hypothesized needs has conditions that add up to resemble our intuitive concept of knowledge closely enough, then it seems we may identify the two. That is, we construct a new concept so that it meets a certain set of standards, and if it has enough features in common with the intuitive concept, then we may describe the intuitive concept's intension as consisting of the new concept's

[10] Contrast Kusch 2009, p. 65.
[11] Craig 1990, p. 14.
[12] Ibid., p. 2.

conditions. Craig finds precedence for such a method of analysis in Rudolf Carnap's theoretical explication.[13] He accordingly refers to his method as a "practical explication", and I follow suit. However, my use of the phrase differs somewhat from Craig's, as I am particularly interested in contrasting the steps described here with the genealogical step that I describe in Chapter Two. Craig simply refers to his entire project, including the genealogy, as a practical explication.

i. Practical explication in four steps

I have broken the practical explication down into four steps:

Step One: Hypothesis

Hypothesize a need that the concept of knowledge might meet.

Step Two: Explication

Analyze the need to determine precisely what would be required in order to meet it. If the hypothesis was well chosen, then the need will require a concept. Explicate the conditions such a concept would require in order to meet the need.

Step Three: Synthesis

Synthesize those conditions into a concept, which, did it exist, would meet the hypothesized need.

Step Four: Comparison

Compare the synthesized concept to knowledge-related intuitions. If the synthesized concept does not contradict the intuitions generated by the concept of knowledge, then that constitutes a strong case for the hypothesis. It would then seem that the concept of knowledge actually does function to meet the hypothesized need. At that point, the burden of proof would be on the opponent.[14]

[13] Ibid., p. 8.
[14] Ibid., p. 4.

ii. Two constraints on the hypothesis

A hypothesis about what the concept of knowledge does for us takes its first form as a hypothesis about our needs. That is, the first step of Craig's method is to ask what kind of need might generate a knowledge-like concept in response.

Craig doesn't find it necessary to provide any arguments for choosing one initial hypothesis over another. Instead he prefers the "proof of the pudding,"[15] meaning that we should evaluate the hypothesis by its results. If the results are good, then the hypothesis is good. That is, if the hypothesized need does in fact demand conditions that resemble our intuitive concept of knowledge, then that constitutes sufficient a posteriori defense that the concept of knowledge does in fact function to meet those needs.

Nonetheless, there are two constraints on the hypothesis from the very beginning.

(1) Generality requirement

The first arises out of the following observation: every society that has a language has a word for the concept of knowledge. Craig takes this to suggest that the concept of knowledge might meet "some very general needs of human life and thought."[16] So, the hypothesized need must be very widely spread. Precisely, it must be common enough to be shared by "any society that has a well-developed language, sufficiently well developed for us to be able to say that it exercises a concept … approximately identifiable with our concept of knowledge…"[17] Thus, the needs must be "so general, indeed, that one cannot imagine their changing whilst anything we can still recognise as social life persists."[18]

If the hypothesized need were more exclusive, peculiar only to certain societies or to certain individuals within any given society, then a practical explication would fail to yield unified results. That is, the analyst would have to construct several concepts, each meeting different needs, in order to hope to describe the different ways in which the concept of knowledge is used. This is a valid alternative, but it is worthwhile to find out if it is our only alternative. That is, we should check whether

[15] Ibid., p. 3.
[16] Ibid., p. 2.
[17] Ibid., p. 4.
[18] Ibid., p. 10.

there might be a widespread need, which one common concept of knowledge functions to meet.

(2) Simplicity preference

The drive to find one unified concept lies in the simplicity preference. Craig relies on the principle that the best explanations are the simplest explanations. Complex hypotheses should be invoked only if simple ones prove insufficient.[19] Thus if the results satisfy, then the hypothesis is complete and ought not to be specified or augmented any further.

The same goes for the results. A more complicated synthesis of the concept of knowledge should only be given if a simpler one is insufficient for meeting the hypothesized needs.

c. The need

i. Survival and information

Now that we have a grasp on how a practical explication works, we may turn to the results Craig draws from it. In accordance with the generality requirement and the simplicity preference, Craig begins with very simple and general truisms about human life.

People need to survive. In order to survive, they must perform survival-conducive actions. In order to perform survival-conducive actions, they must have accurate information about their environment. A human must be cognitively aware of danger in order to best protect himself from it, just as he must be cognitively aware of food in order to best locate it.

> Human beings need true beliefs about their environment, beliefs that can serve to guide their actions to a successful outcome. That being so, they need sources of information that will lead them to believe truths.[20]

These truths serve as premises in humans' practical reasoning. Humans stay alive by avoiding foolish or false reasons for action.

As revealed in the citation above, the need for good information generates the need for good *sources* of information. Often, humans rely on themselves.

[19] Ibid., pp. 4, 68, 92
[20] Ibid., p. 11

> [Human beings] have 'on-board' sources, eyes and ears, powers of reasoning, which give them a primary stock of beliefs.[21]

Craig does not hypothesize that the use of 'on-board' sources generates a widespread need for a knowledge-like concept. I return to this point in section iii below.

ii. Informants

The source of information must often be another person, for any number of reasons. Say that a woman is looking one way, while a man is looking another way; she can double the amount of information at her disposal by asking him what he sees. Of course, spatial differences are not the only causes of informational asymmetry; novices need to rely on experts; the short-sighted need to rely on the far-sighted; scientists need to rely on their colleagues, etc. One person cannot figure everything out for herself; the division of epistemic labor is crucial to survival. As a result, just about every member in every society has needed, needs and will need to find out whom to ask and whom to trust about certain questions. That is, people need to be able to evaluate each other as to whether or not they are good informants. Craig's hypothesis, then, is that people need a concept of a good informant.

Craig moves quickly from the need for good sources of information to the need to *evaluate* sources of information.

> On any issue, some informants will be better than others, more likely to supply a true belief. (Fred, who is up a tree, is more likely to tell me the truth as to the whereabouts of the tiger than Mabel, who is in the cave.) So any community may be presumed to have an interest in evaluating sources of information; and in connection with that interest certain concepts will be in use. The hypothesis I wish to try out is that the concept of knowledge is one of them.[22]

In the cited passage, he relies on the following progression of needs:

(1) The need for information generates a need for good informants.
(2) The need for good informants generates a need to think and talk about evaluating informants.

This transition from (1) to (2) is not entirely obvious. In fact, it relies on a controversial premise about the nature of testimonial interaction, as the next section shall reveal.

[21] Ibid.
[22] Ibid.

iii. Doubt demands explicit trustworthiness

Craig argues that we should expect to find a need for a knowledge-like concept in social, testimonial situations rather than situations in which humans use their 'on-board' sources. That is, the widespread need to attribute knowledge is not based on the need to make first-person attributions, i.e. to attribute an evaluative concept to one's self by saying, "I am a good informant as to whether p." According to Craig, people are more likely to need a way to say something like, "He is a good informant as to whether p."

> To find oneself in possession of a belief on the question whether p pre-empts inquiry; to take a self-conscious look at one's own apparatus with the doubt in mind that it may have delivered a falsehood calls for a considerable degree of sophistication. Our investigation ought to start from the position in which we as yet have no belief about p, want a true belief about it one way or the other, and seek to get it from someone else.[23]

He reasons in this passage that because humans are not so likely to doubt their own beliefs, the need to evaluate one's own trustworthiness might not meet the generality requirement. He does, though, assume that the need to evaluate *others* meets the generality requirement. Craig's hypothesis is thus based on the assumption that *humans are very likely to doubt the assertions of others*. Otherwise, the need for a third-person evaluation would be as unable to meet the generality requirement as the need for a first-person evaluation.

Just as humans only need a concept of *edible* mushroom if there are some *inedible* mushrooms around, humans only need a way to separate the *trustworthy* from the rest if they tend to encounter a good deal of *untrustworthy* assertions. Thanks to the widespread nature of interpersonal doubt, the need for such an explicit evaluation passes the generality requirement. It is the presence of doubt that bridges the gap between the need for good information and the need for a concept of the good informant. People need a concept that describes those situations in which their doubt may be dispelled.

We need not conceive of this doubt as an active suspicion. Even a much weaker form of doubt, say an inarticulate hesitation before committing to a positive belief, could still result in a need for a concept that entails permission to believe. Nonetheless, this may seem to gloss over some controversial points in the

[23] (Ibid.)

epistemology of testimony about whether and how one may, in normal circumstances, take another's assertion to be true.[24] Yet Craig defends the third-person form of the hypothesis by touting its *simplicity*; he claims that first-person considerations would make the synthesis unnecessarily complex.[25] Given the complexity involved in arguing that other people are harder to trust than one's self, I find it unlikely that a restriction to the third-person scenario truly follows from the simplicity preference. Thus, I shall modify his hypothesis to include the needs that arise in first-person scenarios as well. However, for the rest of this chapter, I refrain from offering any critique and simply present Craig's own results, which rely exclusively on the third-person hypothesis.

d. The conditions

The next steps involve explicating the conditions required to meet the need for good informants and synthesizing them into one concept. The inquirer needs to recognize trustworthy subjects. What conditions would a concept have that could meet that need? Craig assumes that the need demands a concept of the good informant. The question then becomes, what features make a subject a good informant? What is the inquirer looking for?

i. Prototypical case

The first interesting point is a negative one: the concept of the good informant need not be definable with a list of necessary and sufficient conditions. In this respect Craig radically departs from conceptual analysis. For such analysts, as Michael DePaul writes, "the primary criterion for the adequacy of a philosophical analysis is immunity to intuitive counterexample."[26] When practically explicating a new concept, though, the only criterion one may presuppose is that the concept must meet the hypothesized need.

Inquirers may need a concept with a certain set of conditions, x, y, and z. Still, it might turn out that, in one unusual situation, x, y, and z are all satisfied, and yet the

[24] For a presentation of these controversies see (Adler, 2010 Edition)
[25] (Craig, 1990, p. 68)
[26] (DePaul, 2009, p. 124)

inquirer would still doubt. There might also be unusual situations in which one of the conditions is not satisfied, yet the inquirer has some other good reason to trust.

> In freakish circumstances, a purpose may be achievable in unusual ways – factors which would usually frustrate it may, if other features of the situation are exceptional, do no damage, factors which are usually vital may, abnormally, be dispensable.[27]

Craig argues that neither of these unusual situations would have the power to change the concept.

I can find at least two motivations for his rejection of the "primary criterion" of conceptual analysis. First, Craig draws attention to how the concept is likely acquired.

> It is precisely by being everyday practice that everyday practice manages to impress itself upon speakers and so stay what it is.[28]

As an inquirer repeatedly searches for and comes across good informants, certain characteristics common to each of these experiences impress themselves upon her. Indeed, any concept that serves a practical purpose for entire communities likely must be acquired in just this manner, for it must be acquirable by every child, not only those with specialized educations. If this is indeed Craig's point, then he would have an ally in Timothy Williamson who supports an analysis of 'knowledge' as a vague term with the following argument:

> We understand [many philosophically significant terms] not by learning precise definitions but by extrapolating from examples which leave their application to ranges of borderline cases unclear.[29]

Language acquisition is a hodgepodge affair and the blurry edges of the concept of knowledge might be expected to reflect that.

Craig also argues that a less determined concept would be better able to play the hypothesized function, as it is better fitted to an inquirer's strategic concerns.

> To try to make a practice of detecting freakish cases would mean incurring high costs in time and energy; and successful detection would scarcely ever offer any benefit which could not be had by finding a standard informant, or investigating for oneself. In practice, therefore, it must be the standard or prototypical case at which the inquirer's strategy is directed, so that one might almost say that for practical

[27] Craig 1990, p. 14.
[28] Ibid., p. 16
[29] Williamson 2005, p. 99.

> purposes what the concept amounts to is the essential description of the prototypical case.[30]

Remember, according to the hypothesis the concept of the good informant exists in order to help inquirers who are looking for an answer to the question "Who can tell me whether *p*?" The concept of the good informant will thus take on those features for which the inquirer is alert when she is trying to find someone who can tell her whether *p*. Her best strategy is to be alert for the features that are *usually* there. If in some bizarre but possible world she might get what she wants without all of those features, or she might have all the features, but still cannot get what she wants, that should not affect her strategy.

This explication from the hypothesized need makes a biconditional definition of the concept of knowledge inappropriate. That is, we will not be able to plug the explicated conditions into the right side of the following formula: "S knows that *p* if and only if x, y, and z." The strictness of the phrase "if and only if" does not properly reflect the vague concept with typical conditions that the hypothesized need requires. Nonetheless, in what follows I will continue to present the results of the synthesis in just such a formula. I do so only to promote the clarity of my claims. The point should not be forgotten that the "if and only if" is not meant to indicate strict necessity and sufficiency. Rather, the conditions on the right side of the formula will *typically* be individually necessary and *typically* be jointly sufficient. Thus, one bizarre counterexample will not be a satisfying refutation of my final definition.

Bizarre counterexamples may be helpful insofar as they *suggest* a serious flaw. While thought experiments might present extremely unlikely situations, they can still inspire very clear and strong intuitions. It is likely that these intuitions do have some relation to the function of the concept of knowledge.[31] Still, opponents need to show that relation; that is, they need to show how the counterexample demonstrates the overall inability of the synthesis to fully account for the function of the concept of knowledge. Therein lays the strength of such an objection. If the counterexample does not reveal any such inability, then it is only a freak curiosity and may be disregarded.

[30] Craig 1990, p. 15.
[31] Ibid., p. 7

ii. The goal: True, intelligible assertion

A demand for the concept of the good informant arises out of the inquirer's need to answer the question, "Who can tell me whether *p*?" Therefore, according to Craig, the conditions of the concept are going to be those conditions that would satisfy the inquirer. What does the inquirer want?

She wants an informant that offers true information as to whether *p*. That offering usually takes the form of a *belief* that the informant *sincerely asserts*. It is possible that the informant offer the information in some other way. Maybe the assertion is true, but only because the informant is lying about his false belief. In that case, so long as the inquirer sees that the informant is lying, she can still get her true information. Maybe the informant has a hesitant constitution and so doesn't quite believe the information. In that case, so long as the inquirer sees that the informant should be more certain, she can still get her true information.

Nonetheless, due to the rarity of such scenarios and the amount of background knowledge they demand of the inquirer, the inquirer's strategy will still be to find informants who will offer information in the form of a true belief that they sincerely assert. Therefore, Craig insists that we may still explicate the condition that the informant must have a true belief as to whether *p*. Thus, if an informant is to be good, then the following condition must hold:

> **Disjunctive belief condition**
> S_1 is a good informant for S_2 that p only if either *p* and the informant believes that *p*, or not-*p* and the informant believes that not-*p*.[32]

Since in the prototypical case the good informant that *p* would believe that *p*, Craig synthesizes the belief condition into his concept. Christoph Kelp, however, remains unconvinced that the good informant concept really would demand a belief condition.[33] I will return to this objection later, offering two further arguments for the explication of a belief condition.[34]

[32] (Ibid., p. 12)
[33] (Kelp, 2011, p. 3)
[34] See section 2.a.iii and 6.c.

In any case, Craig holds that the good informant concept also has an intelligible assertion condition. That is, in order to be a good informant, the informant must sincerely assert his belief to the inquirer, and the inquirer must be able to understand him.[35] We have thus explicated the following conditions for the concept of the good informant:

The truth condition

S_1 is a good informant for S_2 that p only if p.

The intelligible assertion condition

S_1 is a good informant for S_2 that p only if S_2 hears and understands S_1 assert that p.

Notice that I have simplified Craig's version. Rather than keep the disjunctive form, which describes what S_1 needs to be a good informant as to *whether* p, I have only described what S_1 needs to be a good informant *that* p. With this reformulation, we lose the element of uncertainty that characterizes the inquirer's position. However, I think that this element is adequately accounted for by the following indicator property condition; so I will leave the truth and intelligible assertion conditions as they are, for simplicity's sake.

iii. The strategy: Indicator properties

If an inquirer is just looking around for subjects with true beliefs as to whether *p*, she is not going to get very far. How can she separate the subjects with true beliefs from the subjects with false beliefs? She needs some sort of *indication* that the subject has a true belief. That is, an inquirer will not directly search for an informant with a true belief, but rather for an informant with some property indicative of a true belief about whether *p*. To meet that need, Craig holds that we must explicate not only a condition that reflects the inquirer's goal, but also a condition that reflects her strategy. Without a strategy, what use is a goal? Craig's final synthesis includes both. In this aspect, his work coincides with that of Matthew Weiner, who holds that the

[35] Craig 1990, p. 85.

concept of knowledge is, like a Swiss Army Knife, "an economical way to carry its components around."[36]

Craig calls the indicator property simply property X. It needn't be specified any further, for it is not any one particular sort of indicator property the inquirer is after; any sort will do. Therefore, we need not complicate our analysis with stipulations about the informant's belief having been caused by this or that, or about the informant having a good reason. Although both those sorts of properties *could* function as indicator properties, neither can claim to be the *only* desirable indicator property. The only thing that matters is that the inquirer can somehow tell that the informant has got a true belief. To do so, the inquirer must detect any property that correlates reliably with being right about *p*.

The question then arises, *how* reliably must the property so correlate? We may phrase the question of reliability in the terminology of possible worlds. What range of possible worlds must the inquirer take into account when deciding whether the informant is likely to be right as to whether *p*? Craig argues that, here too, the best answer is an indeterminate one. The range of possible worlds in which a good informant must be right as to whether *p* is entirely relative to the inquirer.[37]

Many variables can affect the level of reliability that the inquirer demands. If it is very urgent that the inquirer form a belief, she might accept any informant with a property that indicates better than even chances of being right.[38] If, on the other hand, her stakes are extremely high and the consequences will be devastating if she forms a false belief as to whether p, then she will only accept an informant with an indicator property that correlates extremely reliably with being right about p.[39]

The two clauses on the indicator property are that it (a) correlates with being right about whether p reliably enough for the inquirer and (b) is detectable by the inquirer. We may therefore formulate the indicator property condition as follows:

[36] Weiner 2009, p. 164. Weiner uses this argument to demonstrate that knowledge per se is not valuable, only the *concept* of knowledge is valuable. As discussed in section 1.a, Craig's conceptual synthesis demands nothing more. However, Craig's work does show that the concept's value far exceeds its *convenience* as a compound of strategy and goal, which is the only value Weiner claims for it.
[37] Craig 1990, p. 20.
[38] Ibid., p. 86
[39] Ibid.

The indicator property condition

S_1 is a good informant for S_2 that *p* only if S_2 detects that S_1 has a property X which correlates with being right about *p* reliably enough for S_2's concerns.

iv. The minimal synthesis

Craig stops with these conditions. The concept he has explicated to meet the need of getting true information synthesizes the following conditions: a true belief, an intelligible assertion, and a detectable indicator property.

The concept of the good informant has turned out fairly minimal and vague. Craig argues that it needn't include anything more specific. As the inquirer looks for someone who can tell her whether *p*, she's just looking for someone who truly asserts that *p* and who has some detectable property that correlates reliably with being right about whether *p*.

So the synthesized concept could be characterized as follows, in which the inquirer is S_2 and the informant is S_1:

The Concept of the Good Informant

S_1 is a good informant for S_2 that *p* if and only if (iff) *p*, S_2 believes that *p*, S_2 hears and understands S_1 assert that *p*, and S_2 can detect that S_1 has a property X which correlates with being right about *p* reliably enough for S_2's concerns.

2. Genealogies

a. Objectivization

The concept of the good informant is not equivalent to the concept of knowledge; we cannot take the former as our practical explication of the latter. The two different concepts generate very different intuitions in far too wide and important a range of scenarios. For example, imagine that Jessica wants to find out what time it is. She asks Kashim, but Kashim only speaks Urdu and Jessica only speaks English. Since Jessica cannot understand Kashim's assertions, Kashim is not a good informant for her as to the time. However, we have no corresponding intuition that Kashim does not *know* what time it is. Neither do we intuitively accept that, *relative to Jessica*, Kashim does not know what time it is, for we would not agree that Jessica could truly say such a thing. Indeed, we do not have the intuition that one man's knowledge is relative to any single inquirer. The concept of knowledge is more *independent* from its user's individual situation than the concept of the good informant.

i. Four subjective conditions

Craig recognizes that the concept of the good informant is not equivalent to the concept of knowledge. That is, the following definition would be false:

> S_1 knows relative to S_2 that *p* iff *p*, S_2 hears and understands S_1 assert that *p*, and S_2 can detect that S_1 has a property X which correlates with being right about *p* reliably enough for S_2's concerns.

Intuitively, that definition would be far too restrictive. We do not reserve knowledge attributions for subjects who happen to give us information. Craig traces the discrepancy back to four conditions that must be synthesized as part of the concept of a good informant, but have no place in the concept of knowledge. These conditions arise in response to the inquirer's subjectively varying needs; I will accordingly refer to them as subjective conditions. If I am to recognize a subject as a good informant, I demand that the following conditions be satisfied:

1) He should be accessible to me here and now.
2) He should be recognizable by me as someone likely to be right about p.
3) He should be as likely to be right about p as my concerns require.

4) Channels of communication between him and me should be open.[40]

Intuitions and linguistic behavior both agree that a subject can know something even if he doesn't meet these four subjective conditions. The concept of knowledge cannot be reduced to the concept of the good informant.

However, despite the serious discrepancies revealed in the fourth step, Craig does not simply start over with a new hypothesis. Rather, he makes several modifications to the concept of the good informant. The modifications result in a new concept, which does a better job of matching the intuitive concept of knowledge. Naturally, Craig must avoid saving his hypothesis with ad hoc modifications. He must explain why people would have as strong a functional reason to use the modified concept of the good informant as they would have to use the original concept. He offers such an explanation with the story of the process of objectivization.

ii. Social and diachronic needs

The concept of the good informant refers only to subjects who satisfy the inquirer's own immediate needs. This limited scope is the direct result of Craig's hypothesis. In his prototypical case, an inquirer may need, for example, an answer to the question, "Who can tell me whether that's edible?" because she is hungry and needs to have the information right now. A good informant, accordingly had better be able to tell her right now, or else he cannot meet her needs. As a hypothesis, this sort of need is helpfully vivid, but also extremely restricted. Since the concept of knowledge has a much broader extension, we must conclude that it meets a broader, more general need.

Therefore, to bridge the gap between his results and our intuitions about knowledge, Craig extends his hypothesis, considering further needs. He points out that an inquirer also has diachronic and social needs, that is, needs which concern her position in the *future* and needs which are affected by her participation in a *community* of interdependent subjects. These needs will impact her use of the concept of the good informant. She may need, for example, to find someone who can tell whether that's edible, because she will often be hungry and so might want such information in the future. If a subject would be able to tell whether that's edible, then

[40] Ibid., p. 85

he qualifies as meeting the inquirer's diachronic need even if he happens to be currently indisposed. Likewise, the inquirer needs to notice people who could identify food for others as well. It is not only altruism that makes it worthwhile to pick out people who might be able to help others. As Craig points out, "it isn't necessarily for the sake of the horse that we give it the oats instead of eating them ourselves."[41] A strong community needs strong members; cooperation is necessary to survival.[42] Thus, if the subject would be able to help someone else who was hungry, he meets the inquirer's social need. Thanks to diachronic and social needs, the inquirer is interested in recognizing a subject that *would* be able to tell whether that's edible, if only he were awake and someone were hungry.

Craig describes the relationship between an isolated individual's momentary private needs to a socialized individual's diachronic needs as a *development* which he calls the process of objectivization. The story goes as follows: isolated individuals needed the concept of a good informant. Then, they began to adapt the concept in response to social and diachronic needs. With each adaptation, the concept became more objective, until it was finally the fully objective concept of knowledge we have today.

We could tell the story of objectivization about many other concepts as well. For example, we could tell the following story about the concept of chair. Once upon a time, one man wanted to sit down. To satisfy that desire, he did not need to notice anything besides close and accessible objects.[43] Then he began to think about the desire he may have to sit in the future and about his family's desire to sit. Eventually, he no longer noticed only those objects accessible to him, but anything which could at some point be accessible to anyone. At that point, the fully objectivized concept of chair had been formed and even a folded up unused chair on top of a bonfire could still be included in its extension, so objective had the concept become.

I mentioned above that the indicator property may be understood with the terminology of possible worlds. In order to clarify the modifications wrought by objectivization, I shall apply the same technique here. The subjective conditions

[41] Ibid., p. 83
[42] As Rudyard Kipling puts it in "The Law of the Jungle" from *The Second Jungle Book*, "For the strength of the Pack is the Wolf, and the strength of the Wolf is the Pack."
[43] Craig 1990, p. 84.

emphasize the actual world.[44] Insofar as I am applying the concept of the good informant, I am only considering myself, here and now. Bringing different times and other people into consideration means that I must take possible worlds into account as well. The extension of the objectivized concept will certainly be wider as a result. That is, there are far more subjects that I may recognize would be a good informant for me in some possible world than there are subjects who are good informants for me right now in this actual world. As it happens, at the moment I am not speaking to anyone, much less asking someone for information. So there are, for me, no actual good informants right now. However, there are countless individuals who *could* be good informants for me. When I consider my future needs and the needs of others, I recognize the importance of that untapped potential.

Given this relationship with possible worlds, the concept of the objectively good informant may be described as a counterfactual conditional:

> If certain circumstances were to concur (S_2 needed to find out whether *p* and S_2 was talking to S_1 in a language she could understand and S_2 recognized that S_1 was reliable enough for her), then S_1 would be a good informant for S_2 that *p*.

If that counterfactual is true for all nearby possible worlds, then S_1 may count as an objectively good informant, whether or not the circumstances actually concur in this world.[45] So, the objective concept could end up taking the following form:

> S_1 is an objectively good informant for S_2 that *p* iff in a nearby possible world x, y, and z would hold.

However, we should notice straight away that at least one condition cannot change with objectivization: the truth condition. There is no social diachronic need to accept false information as to whether *p*.[46] Whether one's interest is for oneself or another, for the present or for the future, the interest is always in getting at the truth. S_1 could never be an objectively good informant that *p* if *p* is actually false. Therefore the

[44] With the exception of condition (3). See section a.v below.
[45] Craig 1990, p. 93.
[46] Which is not to say that we could not dream up scenarios in which inquirers would need false information as to whether *p*. However, our confinement to the strategic focus of the prototypical case (1.d) prevents us from allowing such cases to affect the synthesis.

definition above could not be correct, for the following modification is sure to be necessary:

> S_1 is an objectively good informant for S_2 that *p* iff *p*, and in a nearby possible world x, y, and z would hold.

Yet truth is not the only feature people need to recognize in the actual world, regardless of their social and diachronic concerns. The following sections will show which conditions may be satisfied by a nearby possible world and which must be satisfied in the actual world. Two conditions will turn out necessary in the actual world: a belief condition and an indicator property condition. Two conditions will turn out to demand satisfaction only in nearby possible worlds: the intelligible assertion condition and the detectability condition. In fact, given the satisfaction of the former two conditions, the latter two will turn out so easy to satisfy that we will be able to drop them from the synthesis altogether without substantially changing the synthesized concept.

iii. From an assertion condition to a belief condition

The concept of the good informant is modified by the intelligible assertion condition, "S_2 must hear and understand S_1 assert that p," resulting in the subjective conditions (1) and (4). How does objectivization affect the need for intelligible assertion? Craig points out that, in a social context, I care about whether someone will be accessible to others as well as myself and I hope that they will do the same. That is, cooperation makes it worthwhile to notice subjects who could inform others, even if they cannot inform me.[47] As a member of society, I need to recognize whether a subject would be a good informant for *other* inquirers. Furthermore, in a diachronic context, I need to recognize whether someone will be accessible at a later time. So I need to notice subjects who could inform others and myself at some time, even if not now. When the conditions (1) and (4) are objectivized accordingly, they have the following form:

(1) S should be potentially accessible to an inquirer.
(4) Channels of communication between S and an inquirer should potentially be open at some time.

[47] Craig 1990, p. 88.

If we clarify that 'accessibility' means that S is heard, and 'intelligibility' means that S is understood, then we can combine the two objectivized versions to arrive at the following counterfactual:

> If S_1 asserted *p* and S_2 heard and understood S_1, then S_1 would be a good informant for S_2.

When we look for an informant who satisfies that counterfactual, we are looking for an objectively good informant. Using the results of Chapter One, such a concept would be defined as follows (with the modified condition emphasized):

> S_1 is an objectively good informant for S_2 that *p* iff *p*, **S_1 asserts that *p* in a nearby possible world, S_2 hears and understands S_1 assert that *p* in a nearby possible world,** and S_2 can detect that S_1 has a property X which correlates with being right about *p* reliably enough for S_2's concerns.

This definition is clearly incomplete; it leaves us forced to accept a drastic discrepancy with the intuitive concept of knowledge. Imagine that, in the actual world, the truth and indicator property conditions hold, but S_1 asserts not-*p*. Elias is usually a champion bird identifier, but thanks to a random event (perhaps an arrow in the eye) he mistakenly believes that the bird before him is a green finch, when it is actually a linnet bird. Even the most trustworthy can sometimes make mistakes. Given the definition above, this mistaken subject counts as an objectively good informant, for in a nearby possible world, the arrow would not have hit his eye. Thus, in a nearby possible world he would not have made a mistake, but rather correctly asserted that *p*. However, the concerns of objectivization in no way suggest that a *mistaken* inquirer should be considered an objectively good informant. People are interested in identifying subjects who would be good informants if only the obstacles of accessibility and intelligibility were overcome; they are not interested in identifying subjects who would be good informants if only they were accurate. Accuracy is one of the properties that people demand to have satisfied in the actual world.

Yet, the concerns of objectivization do suggest that the subject need not actually be asserting *p* to count as an objectively good informant. How then can we include

the demand for actual accuracy in the concept of the objectively good informant? What is it that must be accurate? I suggest that here we have clear functional support for a belief condition. Remember, it was not immediately obvious that a subject must believe that *p* in order to be a good informant that *p*.[48] All the inquirer actually needs is the subject's true and trustworthy assertion that *p*. Usually, people don't assert true things in a trust-inducing manner unless they also believe it, but still, we can imagine some scenarios in which they might. Craig insists that he must still synthesize a belief condition into the concept of the good informant, arguing that the good informant will almost always believe what he asserts, which makes belief integral to the prototypical case.[49] There is, I maintain, an even stronger demand for the belief condition when we consider the need for an objectively good informant. People judge whether or not a subject *would* assert the truth about *p* by determining whether or not he *actually* believes that *p*. The belief condition prevents the concept of the objectively good informant from including mistaken subjects, who actually believe and actually assert that not-*p* despite the fact that in a close possible world, they would believe and thus assert that *p*.

When looking for an objectively good informant, we look for someone who satisfies the following counterfactual:

> If S_1 asserted his actual belief that *p* and S_n heard and understood his assertion, then S_1 would be a good informant for S_n that *p*.

As long as S_1 does believe that *p*, there will *always* be some nearby possible worlds in which S_2 would hear and understand S_1 assert his belief that *p*. Therefore, without changing the behavior of this condition, we may omit mention of potential assertion altogether and reformulate it as a simple belief condition:

The actual belief condition

S_1 is an objectively good informant that *p* only if S_1 believes that *p*.

[48] As explained in 1.d.ii.
[49] See 1.d.i.

When we add this condition to the entire good informant concept, we get the following synthesis (with the modified condition emphasized):

> S_1 is an objectively good informant for S_2 that *p* iff *p*, **S_1 believes that *p***, and S_2 can detect that S_1 has a property X which correlates with being right about *p* reliably enough for S_2's concerns.

iv. Losing the detectability condition

A subject can know something, despite the fact that no one can detect his indicator property. Craig offers the examples of a secretly studious milkman and a little girl who usually lies, but just now is telling the truth.[50] No one can actually recognize them as good informants. What functional reason could there be that we intuitively agree that they know, despite the fact that their knowledge is useless to others? Craig then points out that in a community, a member may care about whether a subject is a good informant on a certain subject, even if she herself would not be able to detect him.[51] She needs, for example, someone in her community to be a good informant about the mass of propositions necessary to disarm and dismantle a nuclear missile. As long as the right people can detect and understand such a subject, that desire is satisfied; she herself need not be able to detect him.

To account for this feature of social life, we need a modification that resembles the change wrought by substituting the subjective assertion condition for the belief condition. In the last section, I argued that, whether or not S_2 actually *hears and understands* S_1 assert that *p*, as long as S_1 actually *believes* that *p*, S_1 may count as knowing that *p*. Likewise, whether or not someone actually *detects* S_1's property X, as long as S_1 actually *has* some property that correlates reliably with being right about *p*, then S_1 counts as knowing that *p*. The objective detectability condition therefore means inquirers will need to recognize people who satisfy the following counterfactual:

[50] Craig 1990, p. 82.
[51] Ibid. p. 89

If S_2 detected S_1's actual property X which correlates with being right about *p* reliably enough for S_2's concerns, then S_2 would be a good informant for S_2 that *p*.

The range of possible worlds must include all logical possibilities of S_2 detecting the property X without changing S_1's possession of property X. Since, as long as S_1 actually has a property X, there will always be some nearby possible worlds in which S_2 can detect it, the detectability condition can be omitted entirely without changing the extension of the synthesized concept. The following simplified condition results:

The indicator property condition
S_2 is a good informant for S_1 that *p* only if S_1 has a property X which correlates with being right about *p* reliably enough for S_2's concerns.

By modifying the indicator property condition so that it no longer is affected by a subjective detectability condition, we have arrived at the following definition of the objectively good informant (with the modified condition emphasized):

S_1 is an objectively good informant for S_2 that *p* iff *p*, S_1 believes that *p*, and **S_1 has a property X** which correlates with being right about *p* reliably enough for S_2's concerns.

v. The reliability condition

The concept of the good informant has the following indicator property condition:

S_2 must be able to detect that S_1 has a property X which correlates with being right about *p* reliably enough for S_2's concerns.

The indicator property condition has two elements of subjectivity: condition (2), S_2 must be able to detect the indicator property; and condition (3), S_1 must be reliable enough for S_2's concerns. I have just shown how condition (2) may be objectivized and eliminated.

We may call the condition (3) the reliability condition. As discussed in 1.d.iii, this condition demands a consideration of possible worlds even in its subjective version. The inquirer needs to find an informant with a property that correlates with being right about *p* in a range of nearby possible worlds. That range is set relative to the inquirer.

Notice though that the subjective reliability condition does not interact with possible worlds in the same way as the objectivized assertion and detectability conditions. Those latter conditions could be satisfied by any *one* close nearby possible world. The reliability condition, on the other hand, must be satisfied by *all* close nearby possible worlds, otherwise the informant does not count as reliable enough.

Craig argues that the reliability condition must drastically change in a social setting. Members of a society must label good informants without any clue as to who might need their information, or what concerns those indeterminate inquirers will have.[52] Objectivized, the reliability condition therefore reads:

3) He should be as likely to be right about *p* as anyone's potential concerns require.

That is, we widen the amount of relevant possible worlds in which the informant must be likely to be right about *p*. Say that Joel believes that Annie is a redhead. Given our uncertainty about who might need the subject's information and why, we must consider all possible worlds in which he is giving his information. What if, for some reason, the president's decision whether to go to war depended on Annie's hair color – would Joel be a reliable enough informant for the president in that possible world? Would Joel be reliable enough for a shopper deciding which shoes to buy, or for a policewoman catching a criminal? For every possible inquirer with every possible concern, we must consider the possible world in which Joel is telling that inquirer that *p*. If in one single possible world in that drastically widened range Joel fails to be reliable enough for the inquirer using his assertion, then he also fails to be an objectively good informant. Unlike the assertion and detectability conditions, the objectivized reliability condition becomes much more difficult to fill, so we cannot expect it to merely fall away.

[52] Ibid., p. 91

There are a number of important questions here, concerning both Craig's motivation and the resulting modification. Since I answer them differently than Craig does, I'll postpone considering this condition until Chapter Three. Here I shall only foreshadow my further treatment: I do not agree that we need the concept of knowledge to have such an objective reliability condition. I shall accept a version of contextualism according to which the level of reliability demanded by the concept of knowledge is indeed affected by the attributor's concerns. Therefore S_2, qua potential inquirer, will never disappear entirely from the definition of knowledge. The range of possible worlds in which a knower must be reliable does in fact depend in part on the attributor.

vi. The objectivized synthesis

By considering the needs implied by a diachronic social environment and explicating the conditions which could meet them, we arrive at the following synthesis:

> **The Objectivized Concept of the Good Informant**
> S_1 is an objectively good informant for S_2 that *p* iff *p*, S_1 believes that *p*, and S_1 has a property X which correlates with being right about *p* reliably enough for S_2's concerns.

Intelligible assertion and detectability are far less important to this synthesis then they were in that of the concept of the good informant. A subject may count as an objectively good informant as long as, in some nearby possible world, he intelligibly asserts that *p* to an inquirer and, in some nearby possible world, the inquirer can detect his indicator property. I argued that with this modification these conditions become so easy to satisfy that they are dispensable to the synthesis. I also argued that the belief condition gains in importance, so that its place in the synthesis is less disputable.

b. A Taxonomy of genealogies

Craig claims that the concept of the good informant is the origin of the concept of knowledge,[53] and that this origin has left marks.[54] Yet he fails to expound on further on exactly how we should understand these expressions 'origin' and 'mark'. We must clarify the status of the concept of the good informant and its relationship to the new objectivized synthesis. Several attempts to do so have claimed that the concept of the good informant is the first step in the genealogy of knowledge. These attempts have employed varying and in part contradictory presuppositions about how and why to construct such a genealogy. However the differences have not been made explicit; apparent agreement has veiled incompatibility. Thus in the following section, I offer a taxonomy of genealogies, dividing the varieties into four categories: rhetorical, imaginary, conjectural and factual.[55]

i. Rhetorical genealogy: Craig

In order to understand genealogies, first we must consider an aspect of Craig's theory which I have not yet mentioned: his use of the state of nature. His reference to a state of nature places Craig's work in a larger tradition of political and ethical philosophy, practiced by philosophers such as Hume, Hobbes, and Nietzsche. However, despite its prominence in the title, nowhere in his book does Craig explain exactly how we should understand the status of the state of nature. In a later essay[56], he discusses the matter at more length. The state of nature is a tool to be used in a genealogy; its constitution thus depends on the sort of genealogy in which it is used. According to Craig, a genealogy can be factual, imaginary, or conjectural.[57] He seems to think of his own as imaginary; however I prefer to save that term for a rather different approach, discussed below. Therefore, I propose that Craig employs a fourth kind of genealogy: rhetorical.

Craig characterizes his own use of the state of nature as,

[53] Ibid., p. 97.
[54] Ibid., p. 95.
[55] Craig mentions three of these categories in his 2007 essay, "Genealogies and the State of Nature." The precise characterizations which follow, however, are my own.
[56] Craig 2007.
[57] Ibid., p. 185.

> not an argumentative strategy, but a literary device, that of presenting a generalisation about the human condition as a sketchy description of the early life of the race.[58]

Craig writes *as if* knowledge had its origins in a state of nature, where a concept of a good informant sprang up, and *as if* objectivization had occurred throughout time, gradually modifying that concept into the concept of knowledge we have today. He could have just as well written without that literary device and omitted the state of nature altogether, without altering his hypothesis or his results. The question then remains what we should make of the process of objectivization. If the genealogy is rhetorical and not meant to be taken literally, then what is objectivization supposed to evoke? If the process of objectivization is not a literal process in time, then how should we understand the modifications? Craig does not address this question, so I offer my own suggestion below (2.c.iii).

On this reading, the state of nature is, as Craig puts it, a "non-load bearing frill rendered harmless by the basic character of the particular facts in question."[59] It involves no actual historical claims, but is rather meant to illustrate very general features of humanity. The features should be general enough to be plausibly attributed to prehistorical societies, but most importantly, they must still be attributable to contemporary societies. In addition to his own work, Craig describes Hobbes' use of the state of nature as also falling under this ahistorical model.[60]

Why rely on this particular rhetorical device? By describing his hypothesized needs as present in a state of nature, Craig ensures generality, naturalism and simplicity. The state of nature is able to help us with these three points for two reasons. One, it is abstract. Two, insofar as it is not abstract but rather reminiscent of a prehistoric tribe, it is very different from our own society.

By bringing our attention to the epistemic needs of a prehistoric tribe, Craig ensures that the hypothesized needs will meet the generality requirement. Remember, he began with the observation that any society with a language has a word for 'knowledge'. Therefore the needs to which the concept of knowledge functions in response must be shared by any society with a language. We prevent ourselves from narrowing the applicability of our practical explication by asking straightaway

[58] Ibid., p. 193.
[59] Ibid.
[60] Ibid., pp. 192 – 193

whether our hypothesized needs could also be attributed to a society very different from our own.

Second, the state of nature is helpful because it ensures naturalism.[61] Of course, prehistoric tribes probably had lively cultures. However, when discussing the needs of a "state-of-nature society" we omit any mention of idiosyncratic cultural influences and focus only on general *practical* needs. As a result, the concept we explicate as a natural response to those needs will be a natural response to the humans' interaction with their environment. We thereby avoid any inadvertent reversion to Platonic presuppositions.

In a related way, the state of nature can help to ensure simplicity. With an abstract model of a society as our starting point, we avoid considering any elements extraneous to the basic natural needs that can get the explication going.

In Craig's work, the "non-load bearing frill" is a rhetorical device, used for the advantages outlined above. Thus, I have characterized his use of the state of nature, and the temporal terms in which he presents the process of objectivization, as a rhetorical genealogy. The other three forms of genealogy are more substantial. I discuss below how Williams and Fricker use an imaginary genealogy, while Kusch and Kelp use a conjectural genealogy. Both Williams and Kusch stress the additional importance of a factual genealogy.

ii. Imaginary genealogy: Williams

The term 'genealogy' was introduced in relation to Craig's work by Bernard Williams in his book *Truth and Truthfulness.*[62] Williams goes on to offer a genealogy of his own, not of knowledge, but rather, of the practice of truthfulness. Williams' genealogy has nonetheless repeatedly been linked to Craig's,[63] so it is worth presenting here in some detail.

Williams conceives of imaginary genealogies as fictional stories. The state of nature and the supposed development out of it are made-up; they are not meant to represent any real society or historical progression. He defends the use of an imaginary genealogy as a way to describe a concept or practice as functional without reducing it to its functional role.

[61] Craig 1990, p. 9. See also Williams B. 2002, pp. 22 – 27.
[62] Williams B. 2002, p. 21.
[63] Craig 2007; Fricker 2007; Kusch 2009; Kusch forthcoming.

> Why not just give a functional account without the story? Do the diachronic fictions of genealogy add anything expect colour? They do. In relation to institutions, practices, expectations, and values that actually exist, of justice, promise-keeping, truthfulness, and so on, functional accounts are simply false. …it is just not true that the dispositions of truthfulness that we have, or that anyone else has had, can be adequately explained in functional terms. …their value always and necessarily goes beyond their function. Nevertheless, at a more abstract level, function plays a role in explaining them.[64]

Williams uses an imaginary genealogy of truthfulness to avoid reducing the value of truthfulness to its merely functional value.

Like Craig, Williams begins with a state of nature in which people need to pool information.[65] He describes how the inhabitants of this state cannot do so if they cannot trust one another to tell the truth. As a result, they begin to value accuracy and sincerity.[66] That is, the individuals begin to be accurate and sincere to ensure that they will all get the true information they need to survive.

This little story tells how the need to pool information caused the tendency to value accuracy and sincerity. The causal element (*"As a result…"*) constitutes the fictional aspect of Williams' genealogy. The factual moral we may draw from the story is that truthfulness does offer the benefits of true information and thereby survival. We judge the moral as true because we can judge the inhabitant's decision as rational.[67] "Of course they needed to start being truthful to one another in order to reap the benefits of true information and survival!" we may think, upon reading the story.

Williams goes on to point out that these benefits cannot suffice to account for the value of truthfulness.[68] For the practice of truthfulness to actually result in the benefits described, people must practice it precisely at those moments in which they themselves are not benefitting. Yet, if truthfulness had only instrumental extrinsic value, there would be no reason for them to sacrifice their own personal gain for the sake of being truthful. That is, if truthfulness were only valuable insofar as it brought individual benefits, then people would lie and be lazy (and hence inaccurate) whenever it suited them, which would likely be often. As a result, people would not

[64] Williams B. 2002, pp. 34 – 35.
[65] Ibid., p. 43.
[66] Ibid., p. 44.
[67] Ibid., p. 34.
[68] Ibid., p. 85.

be able to pool information anymore. If truthfulness only had instrumental value, then in fact it would not have any value at all.

However, people are in general truthful. That is, many, if not most, people continue to be truthful even in situations in which it does not directly benefit them. The cause of this behavior is that people value truthfulness *intrinsically*.[69] They take accuracy and sincerity as being good per se, regardless of whether it benefits them personally. Thanks to this tendency to value truthfulness intrinsically, truthfulness does in fact have instrumental value as well. The wide-spread tendency people have to intrinsically value truthfulness creates communities in which people can pool information and thereby survive. Williams stresses that truthfulness can only have this extrinsic value so long as it is also intrinsically valued.[70]

Thus, Williams makes two points: first, he claims that we must be truthful, or else we will lose the benefits of pooling information. Second, he claims that the social bolsters of the intrinsic value of truthfulness are not merely a façade, *disguising* the actual extrinsic value of truthfulness as a means to the benefits of getting true information. Rather, the social bolsters of the intrinsic value of truthfulness *enable* truthfulness to provide those extrinsic benefits. It is difficult to make both of these points without seeming to assert a contradiction. If he used real examples to support the first claim, he might seem to be denying the second claim. Therefore, he supports the first claim by telling an explicitly fictional story in which people decide that they need to be truthful in order to reap some desirable goods. We grasp the story's moral by recognizing the rationality of their decision. Williams can then base his argument that truthfulness is a socially functional practice which should be established on the moral of that story without falsely implying that real individuals actually should or could use truthfulness as a functional means to an end.

In sum, Williams uses an imaginary genealogy to offer a non-reductive description of the relationship between a practice's extrinsically valuable functional aspect and its historical manifestation as an intrinsic value.

Factual genealogies also play a large role in Williams' work. He points out that the fiction can only explain so much.[71] His imaginary genealogy does not show us how societies actually encourage people to value truthfulness. To understand how

[69] Ibid., p. 92.
[70] Ibid., p. 59.
[71] Ibid., p. 39.

accuracy and sincerity have been socially established as intrinsic values, we need real history.

iii. Imaginary genealogy: Fricker

Miranda Fricker takes Williams' development of imaginary genealogies and applies it to Craig's epistemological work.[72] Like Williams, she constructs an imaginary genealogy that begins with a simplified state of nature in which people need to get information from one another and ends with the idiosyncratic epistemic practices of actual societies. However, Fricker does not use Williams' defense of imaginary genealogies. That is, she does not argue that the genealogy is a useful tool for explaining a phenomenon's functional value without reducing its value to its function. Indeed, she argues that Williams should have written the intrinsic value of truthfulness into the state of nature.[73]

While such a criticism does not make sense given Williams' own motivation for his imaginary genealogy, Fricker's criticism does make sense from her own perspective. She uses imaginary genealogies towards a different end than Williams does and thus operates different standards for how the state of nature should be constructed. Namely, Fricker uses the state of nature and the imaginary genealogy to distinguish a phenomenon's *necessary* features from its *contingent* features.

> A good genealogical explanation of the concept of knowledge helps us understand how, and in what respects, our actual epistemic practices are the contingent social manifestations of our most basic epistemic predicament. So it helps us understand to what extent features of our actual practice are necessary, and to what extent they are contingent. This will in turn explain how some kinds of criticism of our practice are worth making, and how some are senseless. (In particular, it explains why some kinds of political criticism of the norms surrounding rational authority are worth making, and why others can never be genuinely political: where the norm in question is necessary, political criticism is at best futile.)[74]

As the quoted passage demonstrates, Fricker's political goals motivate her use of the genealogy. Necessary features are important to distinguish from contingent features, because productive political critique must restrict itself to the latter.

Fricker portrays Craig's work as an imaginary genealogy so that she may emphasize that people *necessarily* look for indicator properties that correlate reliably

[72] Fricker 1998; Fricker 2007.
[73] Fricker 2007, p. 112.
[74] Fricker 1998, p. 165.

with being right about a given proposition. Without that feature, no social epistemic practice could thrive and survival would be unlikely. On the other hand, which indicator properties a society actually acknowledges and whether they are actually reliable are *contingent* questions.[75]

Fricker uses this Craig-based imaginary genealogy to argue that societies necessarily acknowledge certain indicator properties based on prejudice and fear of outsiders, rather than the properties' actual reliability.[76] In other words, epistemic injustice is present in the state of nature, as is a complementary demand for the virtue of epistemic justice. Thus, it would be unreasonable to attempt to root out epistemic injustice; for it is a necessary result of social life. One must, rather, attempt to train one's sensibility to notice and lessen its effects.[77] I shall not explore these claims further; I mention them only to clarify the motivation behind Fricker's use of the imaginary genealogy as a tool to distinguish necessary from contingent features.

iv. Conjectural genealogy: Kusch

In his pair of essays on Craig, Martin Kusch claims that the concept of the good informant is an ancestor of the concept of knowledge, which he refers to as protoknowledge.[78] While he agrees that knowledge developed out of protoknowledge through a process of objectivization, he also modifies and expands upon Craig's work. In so doing, he draws on Williams' genealogy of truthfulness. I shall discuss three crucial points in which Kusch's treatment of Craig resembles Williams' work.

Williams describes the difference between the extrinsic value of truthfulness as a means to true beliefs and the intrinsic value of truthfulness for its own sake. Kusch generalizes from Williams' work to make the broader claim that not only truthfulness, but the entire institution of testimony has intrinsic value.[79] He uses this claim in conjunction with the testimonial nature of protoknowledge to respond to the swamping problem. The swamping problem demands an explanation of why knowledge has more value than mere true belief. Kusch presents a solution by distinguishing between the extrinsic value of a mere true belief and the intrinsic

[75] Ibid., p. 168.
[76] Fricker 2007, pp. 117 – 118.
[77] Ibid., pp. 81 – 97.
[78] Kusch 2009; Kusch 2011.
[79] Kusch 2009, pp. 75 – 76.

value of a true belief produced by testimony (protoknowledge).[80] This suggestion merits further attention, but it is not directly related Craig's genealogical story of objectivization. Objectivization does not account for the difference between a mere true belief and protoknowledge, but rather the difference between protoknowledge and knowledge. That is, objectivization tells a story that links a true belief in an isolated testimonial interaction (protoknowledge) with a true belief in a diachronic social environment (knowledge). As my momentary aim is concerned exclusively with the latter relationship between protoknowledge and knowledge, I will focus instead on the second two similarities with Williams.

Kusch's treatment of objectivization draws on Williams' work on truthfulness in two ways. First, Kusch praises Williams' use of factual genealogy, stressing the importance of the relationship between a society's set-up and its understanding of knowledge.

> One of the most valuable aspects of genealogy is its systematic use of the idea that the evolution of concepts and the development of social relations are inseparable. Every step in the evolution from protoknowledge to knowledge is explicated in terms of changed needs of the group or changed forms of interaction.[81]

Second, Kusch agrees with Williams that a genealogy is a fiction.[82] So, the development from protoknowledge into knowledge is not meant to be taken literally. Protoknowledge might not have existed at all.

However, Kusch's use of fiction is not compatible with Williams' use of fiction. Recall that Williams justified his use of a fiction with the claim that a reductively utilitarian conception of truthfulness could never be actually held by the members of any society. Were an actual group of people to view truthfulness as instrumentally useful, then truthfulness would never gain any footing at all. The fiction of the state of nature is not meant to be applicable to any single, early society. Rather, the fiction of the state of nature is meant to illuminate the functional value of the practice of truthfulness shared by every single actual society. Williams argues that state-of-nature truthfulness, 'prototruthfulness' if you will, is not an ancestor of real historical

[80] Ibid., p. 76.
[81] Ibid., p. 70.
[82] Ibid., p. 65.

truthfulness but rather a "necessary, structural feature" of real historical truthfulness.[83]

Kusch's position is incompatible with a Williams-style imaginary genealogy, for he argues that protoknowledge is not a feature of knowledge at all. He claims rather that protoknowledge *developed into* knowledge. Furthermore, he claims that the conceptual development was a response to social development. That is, as demonstrated by the quoted passage above, Kusch applies a historical orientation to protoknowledge as well as to knowledge. Accordingly, his genealogy demands that it be possible for there to have been, at one point in real history, a society that used the concept of protoknowledge. He conjectures that protoknowledge did at one point exist, likening the relationship between protoknowledge and knowledge to the relationship between *homo erectus* and *homo sapiens.*[84] Clearly, on this model the concept of the good informant is not in any sense a feature of the concept of knowledge, but rather its predecessor. As we saw, Kusch does make use of Williams' description of intrinsic value. However, in that discussion, the relationship described is between protoknowledge and mere true belief, not protoknowledge and knowledge. Therefore, I suggest that Kusch's work on protoknowledge should be classified as a *conjectural* genealogy, rather than an imaginary genealogy. He offers a conjecture as to how the actual concept of knowledge might plausibly have originated. If the conjecture is false, then the genealogy is mistaken. In that respect, a conjectural genealogy differs from an imaginary genealogy, in which the described 'origin' may very well be false without endangering its relevance.

Kusch argues further that the genealogical story has relevance today in an observable "pattern of use"; people regularly treat the concept of 'knowledge' in ways analogous to how they would treat a concept of the good informant.[85] While we obviously cannot claim that *all* knowledge attributions fall under that very subjective concept, he suggests that knowledge need not be unified under one definition or one function. Rather, we should draw on Wittgenstein's model of a family resemblance.[86] Wittgenstein claims that certain concepts, such as 'game,' can never be adequately captured by one definition. Rather, individual examples fall under the concept thanks

[83] Williams B. 2002, p. 93.
[84] Kusch 2011.
[85] Kusch 2011, pp. 23 – 24; Kusch 2009, pp. 80 – 82.
[86] Wittgenstein 2006, p. 278.

to various similarities with other examples. Staring contests resemble poker in that both have a winner and a loser; poker resembles solitaire in that both use playing cards. There is no one-size-fits-all account that could lay down individually necessary and jointly sufficient conditions for the concept of game.

According to Kusch, the concept of knowledge is also such a concept. Until now, epistemologists have only avoided recognizing that fact by arbitrarily legislating the ways in which people actually use the concept. With no adequate justification, they divide that data into three groups: the 'core' use, which interests them and for which they construct their semantic theory; a pragmatic use; and mere errors.[87] Despite his rejection of strict analytic definitions, Craig too falls under Kusch's critique, since he wrote that the concept of the good informant is the 'core' of the concept of knowledge.[88] If knowledge is a family resemblance concept, then it has no core. Sometimes a certain set of conditions will govern our knowledge attributions, while other times those same conditions will play no part. Specifically, Kusch argues that some knowledge attributions function to meet the needs hypothesized in Craig's practical explication, whereas others do not.[89]

The hypothesis that the concept of the good informant appears as a pattern of use may be supported by a conjectural genealogy.[90] Nonetheless, it does not require it. After all, one may characterize a limited pattern of use of the word 'knowledge' as serving a certain function, without also arguing that the concept of knowledge originated in order to play that function. The pattern-of-use hypothesis may not be supported by an imaginary genealogy, at least, not along the model of Williams or Fricker. Both of their imaginary genealogies describe the elements of the state of nature as special features of the target practice. These features demand to be explicated in an imaginary genealogy in order to illuminate that specialness: for Williams, their functional value; for Fricker, their necessity. If protoknowledge developed into a distinct 'pattern of use,' then an imaginary genealogy would show the special role that the protoknowledge feature plays in the overall pattern of use. However, on Kusch's reading, the pattern of use that interests him is *identifiable* with protoknowledge. There is no room for an imaginary development because there are

[87] Kusch 2011, p. 19.
[88] Craig 2007, p. 191.
[89] Kusch 2011, pp. 23 – 24.
[90] Ibid., p. 22.

no further features to add. If an imaginary genealogy were to support the pattern-of-use hypothesis, it would require a different motivation and a different structure than Williams' or Fricker's.

In sum, we may understand the concept of the good informant in five ways.

(1) As a merely rhetorical device, handy for making some points, but entirely dispensable. Craig characterizes himself along these lines.

(2) As the functional feature of the actual concept of knowledge. This interpretation would be along the lines of Williams' imaginary genealogy.

(3) As a necessary feature of the actual concept of knowledge. Fricker offers this interpretation.

(4) As a coherent concept in its own right, which may have been used at some point early in human civilization and which may have developed into the actual concept of knowledge. Kusch's conjectural genealogy suggests this reading.

(5) As a pattern of use of the actual concept of knowledge, connected to other patterns by family resemblance. This aspect of Kusch's work is compatible with the rhetorical and conjectural genealogies, but not with the imaginary genealogy.

Finally, there is the factual genealogy, the study of how the concept of knowledge has actually been used throughout history. Williams, Fricker and Kusch all accept the importance of such work and engage in it to a certain extent. I see no troubles of compatibility to prevent it being a complement to any of the above approaches.

v. Conjectural genealogy: Kelp

Christoph Kelp follows Kusch's use of a conjectural genealogy; that is, he offers a hypothesis about the conceptual needs of our early ancestors and then shows how changing social needs would bring about changes in the concept.[91] Kelp however, rejects Craig's hypothesis. While Kelp acknowledges that our ancestors must have needed to evaluate informants, he points out that they must have had other conceptual needs as well. Namely, they also must have needed to "evaluate various

[91] Kelp 2011, p. 2.

inquiries agents undertake."[92] Accordingly, Kelp's alternative hypothesis runs as follows:

> ...the concept of knowledge answers this need: it flags when agents *may (permissibly) terminate inquiry* into a given question.[93]

He imagines a group sending out a member to inquire, and asks what properties they would want him to have upon terminating his inquiry. From this hypothesis, Kelp explicates a belief condition, a truth condition, and the following epistemic condition:

> The inquirer's belief on whether *P* stems from a source that is sufficiently trustworthy on the question whether *P*.[94]

He goes on to modify the mismatches between the concept of permissible inquiry termination and the concept of knowledge by accepting both Craig's story of objectivization and Kusch's interpretation of it as a historical development in response to ever more complex societies.

Kelp finds that his account has a crucial advantage over Craig's. The concept of the good informant arises as a need for true *assertions*. Although, on my account, the objectivized version substitutes a belief condition for an assertion condition, this would likely not satisfy Kelp. I imagine he might object to this substitution on the following grounds: why should we narrow the range of possible worlds to only those in which the informant believes that *p*? There is no reason for an inquirer to ever worry about an informant's personal beliefs. We need only narrow down the possible worlds to those in which the informant's *indicator property* does not change. As long as the informant would, in all of *those* possible worlds, always assert that *p*, then he should also count as an objectively good informant. So, despite the forces of objectivization, the concept of the good informant will always extend to include subjects with two kinds of reliability: reliability as a result of the process of belief formation *as well as* reliability as a result of the process that leads from belief to assertion.[95] The concept of knowledge, on the other hand, is only concerned with the processes of belief formation.

---

[92] Ibid., p. 11.
[93] Ibid.
[94] Ibid.
[95] Ibid., p. 9.

Craig does have a way to respond to this objection, which I have not yet discussed; namely, informants who do not believe what they are saying are not properly informants at all, but rather mere sources of information. I discuss this distinction further in section 6.c.

However, no matter what possible defenses are available to Craig, we must take any alternative hypothesis seriously. If an equally plausible hypothesis leads to a different synthesis that nonetheless matches up equally well with our intuitions about knowledge, then the method of conceptual synthesis offers us no tools with which to reasonably judge between the two. Luckily for Craig, Klemens Kappel demonstrates that Kelp's alternative hypothesis does not lead to contradictory results at all.

c. Conceptual synthesis without a genealogy

i. Purified practical explication: Kappel

If a practical explication fails to achieve a satisfying comparison in step four, we need not necessarily turn to a genealogy to bridge the gap. We may rather modify the hypothesis so that the final synthesis does match up with intuitions about knowledge, removing the need for a genealogy entirely. Klemens Kappel takes this route.[96] He argues that the concept of the good informant doesn't match the concept of knowledge simply because Craig begins with the wrong hypothesis. Thus, Kappel practices conceptual synthesis without resorting to any genealogical modifications. Since he thereby avoids the complications and confusions involved in making genealogical claims, he calls his approach "a purified version of what Craig actually does in his book."[97]

Rather than imagine what needs our ancestors may have had, Kappel begins by asking what needs people in general have – our current selves included.[98] He hypothesizes two such needs: the need for an inquiry-stopper and the need for pooling information. That is, he acknowledges the plausibility of both Kelp's and Craig's hypotheses. However, Kappel does not motivate the need for an inquiry stopper by considering situations in which a group sends out a member to inquire.

[96] Kappel 2010.
[97] Ibid., p. 25.
[98] Ibid., pp. 6 – 7.

Rather, he articulates the need in the first-person.[99] That is, individual subjects need to be able to decide for *themselves* when inquiry may permissibly be stopped.

> We need a way to command a switch of attention away from further uneliminated non-p possibilities, a way to urge that we simply take the truth of p for granted in our practical deliberation, as well as in our inquiry into other questions. Or, as I shall say, we have a need for an inquiry-stopper.[100]

People need certainty, but they need expediency as well. It would often be foolish for people to spend so long double- and triple-checking one fact that they miss out on five new facts. Thus, people need a way to say, "that is enough now! Moving on… " Kappel explicates the following predicate which would express the judgment that would meet that need:

> p, and S is in a sufficiently good epistemic position with respect to p, such that S ought to take the truth of p for granted in her practical and theoretical deliberation.[101]

In addition, Kappel accepts the plausibility of Craig's hypothesis as well: people need to pool information and they need to evaluate each other's trustworthiness in order to do so. Thus, Kappel also explicates a predicate which would express the judgment that would meet Craig's hypothesized need, i.e. the judgment that the informant was sufficiently trustworthy. It turns out to bear a great resemblance to the inquiry-stopper predicate:

> p, and S_1 is in a sufficiently good epistemic position such that, given the right circumstances of transmission, S_2 ought to take the truth of p for granted in her practical and theoretical deliberation.[102]

Although the two hypotheses may seem different to begin with, the similarity between the two predicates show that they are actually two instances of the same, more general, conceptual need. Indeed, Craig acknowledges this, although he speaks of the latter as the first-person application of the former.[103] As we have seen above in section 1.c, he assumes that only the third-person version satisfies the generality requirement and therefore prefers to formulate his final synthesis in third-person terms.

---

[99] Kelp also briefly acknowledges this alternative motivation for his hypothesis in Kelp 2011, p. 11.
[100] Kappel 2010, p. 76.
[101] Ibid., p. 78.
[102] Ibid., p. 79.
[103] Craig 1990, p. 67.

Kappel takes a different approach. He combines the two predicates, offering the hybrid predicate as the result of (step three of) his practical explication:

> K(S_1, S_1-S_n, p) iff p, and S_1 is in a sufficiently good epistemic position such that S_1-S_n, given right circumstances of transmission, ought to take the truth of p for granted in their practical and theoretical deliberation.[104]

Rather than assert that either the first- or third-person perspective is in some sense primary, Kappel's K-predicate can describe both scenarios equally well. When an isolated inquirer needs an inquiry-stopper, she would use the predicate to refer to herself by substituting "S_1 – S_n" with "S_1". When an inquirer needs a good informant, she would use the predicate to refer to someone else by substituting "S_1 – S_n" with S_2. I return to the implications of this perspective change for the internalism/externalism debate in the following section.

Kappel does not go on to carry out step four; he does not argue that the "K-Predicate" is in fact equivalent to "knowledge." He only mentions that he personally does believe that the concept of knowledge serves to meet his hypothesized needs.[105] He claims not to need a genealogy or a story of objectivization, because the needs he begins with are contemporary. So, his practical explication is not meant to result in an *ancestor* of the concept of knowledge, nor a *feature* of the concept of knowledge, but in an explication of the concept of knowledge itself.

Kappel does not imply that knowledge attributions *always* function as a K-predicate. That is, it would be entirely possible to read him as claiming no more than that knowledge attributions *sometimes* serve this function. To that extent, his work could be compatible with Kusch's pattern-of-use hypothesis according to which the practical explication illuminates a pattern in which knowledge attributions are used. As I claimed above, one may coherently adopt this scope for one's conceptual synthesis without offering any conjectural genealogy.

Yet even with this weaker claim, Kappel's hypothesis might run into trouble in step four, when we compare his K-predicate to the intuitive concept of knowledge. Notice that a subject deserves the K-predicate as long as he is in a good enough epistemic position *for the given inquirer* and the needs of her deliberation. Yet many inquirers may have extremely lax requirements, while other inquirers have extremely

[104] Kappel 2010, p. 79.
[105] Ibid., p. 80.

stringent requirements. Imagine that Ida and Estelle both ask Nelville whether a given batch of cookies has peanuts. He says it doesn't, but he isn't *absolutely* sure. Now imagine that Ida wants to know because her children are deathly allergic to peanuts, whereas Estelle wants to know because she dislikes the taste of peanuts. Given the differences between Ida and Estelle, Nelville deserves the K-predicate in relation to Estelle, but not in relation to Ida.

The moral of this example is that, without a story of objectivization at his disposal, Kappel's K-predicate allows for a great deal of variability in how reliable a subject must be in order to count as knowing. Is this variability a discrepancy between the K-predicate and the concept of knowledge? That depends on whether or not the concept of knowledge is variable too. As we will see in Chapter Three, there is a school of thought which claims to have empirically found just such a variance in the linguistic use of the word 'knowledge', namely, *contextualism*. That is, the validity of Kappel's hypothesis requires the validity of contextualism about knowledge. Of course, contextualism about knowledge is controversial; one cannot simply assume its validity without further defense. In the following chapters, I provide such a defense, but show how I must refine Kappel's formulation in order to do so.

ii. A contextualist response to the internalist/externalist debate

In what follows, I imitate Kappel's neutrality as to whether the concept of knowledge functions in the first-person or the third-person, assuming that it is flexible enough to do both. Where does this leave me with the question of whether an internalist or an externalist account of justification is correct? Internalists assume that justification entails an element of self-awareness; a subject cannot have a justified belief if she is utterly unaware of the grounds that justify it. Externalists portray justification as possible without that extra element. A reliably produced belief is justified, even if the subject is not aware of that reliable production.

In Craigian terms, the question becomes: must a subject be aware of her own indicator property in order to count as knowing? As far as Craig's original hypothesis goes, the internalist state of self-awareness is obviously not necessary. For an *informant* to count as knowing, the *inquirer* must be able to detect his indicator property. It is unimportant to the inquirer whether the informant can detect his own

indicator property as well. As noted above, Craig does acknowledge that an inquirer can apply the concept to herself, asking, "Am I a good informant as to whether *p*?"[106] To answer that question in the affirmative, one must be aware of one's own indicator property. Thus, one must have the internalist justification, which includes the extra condition of self-awareness. However, as Craig puts it, "the fact that we are in that extra state [of awareness that one possesses X] doesn't oblige us to build it into the concept."[107] Since he portrays the third-person good informant concept as the 'core' of the concept of knowledge, it would be in defiance of the simplicity preference to include conditions demanded by a first-person application in his final synthesis.[108] Thus, Craig's final synthesis is best characterized as offering an externalist account of justification.

Kappel's formulation does not rely on a *core* function. He assumes that the concept of knowledge functions both as a label of good informants and as a label of inquiry stoppers; neither role is somehow primary. Thus, his final synthesis is neutral as to whether the concept is applied in the third- or the first-person. That is, whether "$S_1 - S_n$" is substituted by "S_1" (a first-person attribution) or "S_2" (a third-person attribution) depends on the context in which it is used. Therefore, the answer to the internalist/externalist debate is simply: it depends.[109] When attributing knowledge to a subject, the attributor must have some hypothetical relationship in mind. If the attributor is considering the subject's hypothetical relationship with herself, then the subject must have internalist justification. If the attributor is considering the subject's hypothetical relationship with another, then externalist justification could suffice. Thus, depending on the context of the attributor, either internalism or externalism holds. If this synthesis is correct, it is no surprise that the internalism/externalism debate generates conflicting intuitions: the concept of knowledge itself offers no decisive verdict. [110]

[106] Craig 1990, pp. 63 – 64.
[107] Ibid., p. 65.
[108] Ibid., p. 68. For some doubts concerning his attitude towards the simplicity of restricting the core hypothesis to the third-person scenario, see section 1.c.iii.
[109] Craig shows the similarity between demanding an internalist condition and demanding that the subject know that she knows. (Ibid., pp. 65 – 66) For an analysis of how the demand for the knowledge of knowledge is sensitive to context, see Williamson forthcoming.
[110] For an introduction to the debate and different interpretations of its intractability, see Steup 2008.

iii. Objectivization without a genealogy

We may honor Craig's insights concerning objectivization without using a conjectural or an imaginary genealogy. Craig's practical explication begins with the hypothesis that people need to identify good informants. By explicating the relevant conditions, he synthesizes the concept of the good informant. He then demonstrates how objectivization would modify the relevant conditions. Yet, rather than portray objectivization as an imaginary modification of the explicated conditions, we can understand objectivization as a *correction of the original hypothesis*. In this way, we can incorporate the points Craig makes about objectivization without giving the genealogical aspect anything more than a rhetorical reading.

Craig's original hypothesis only considered the needs of isolated inquirers at isolated points in time. While it did require that these isolated inquirers encounter potential informants, the meeting was artificial. It was artificial, because Craig did not consider how the inquirers' conceptual needs are affected by their membership in a diachronic and social context. Craig's state of nature is a snapshot of an inquirer and informant meeting in a spotlight on a dark stage.

This oversight is best understood as a deliberate oversimplification, not a suggestion that the subjective concept of the good informant might really have once existed. We may accordingly take the considerations of objectivization as a reminder of how society and time affect individual conceptual needs. While it might be easier for us to begin by imagining the instantaneous needs of an isolated individual, we should not consider our hypothesis complete until we have imagined the enduring needs of an individual in a society. Only then do we recognize the importance of concerns for the future and of group action, both of which weaken the accessibility condition; and of the division of epistemic labor, which weakens the detectability condition. If objectivization is to be taken as a conjectural genealogical development, then we must conjecture that early human societies did not worry about the future, were not involved in group action, and did not divide epistemic labor. I find it unlikely that a society in this condition could still nurture the rudimentary testimonial practice necessary for the concept of the subjectively good informant. It is indeed reasonable to conjecture that worries about the future, group actions, and the division of epistemic labor have manifested *in different ways* for different societies at

different times. In more modern societies it could be that these practices are more common and obligatory, more stably established, and that they take place on a larger and more intricate scale. Nonetheless, these are variations on a theme, a theme which is captured by the general conditions on the objectivized good informant concept. Any society with a practice of testimony will need the objectivized concept of the good informant; that need meets the generality requirement. As long as the hypothesis does meet the generality requirement, we will not need to conjecture that the general need has changed throughout time; we need only show how meeting that same need has demanded different sorts of means throughout time.

Therefore, I do not agree that we should conceive of objectivization as a conjectural genealogical development from the need for a good informant to the need for an objectively good informant. So, whereas Kusch describes the relationship between the subjectively good informant and the objectively good informant as related to "changes in patterns of interaction in the group,"[111] it would be more accurate to say that the former focuses on one artificially isolated interaction, while the latter fully accounts for the relationship between group interactions and conceptual needs. We should begin by hypothesizing the more objective, social needs in the first place. By beginning with social needs, we explicate objective conditions.

Due to my diagnosis of the importance of objectivization, in this thesis I synthesize the concept of knowledge by following the four steps of the practical explication and make no recourse to genealogical modifications. My defense is the same as Kappel's: I hold that, when we use my hypothesis, no genealogical modifications are necessary for the explicated concept to resemble the concept of knowledge. My hypothesis is already objectivized; it presupposes a diachronic social environment, so the conditions I will explicate from it will be objective enough. I will leave it an open question whether or not the concept of knowledge might have originated to play the function I describe. I claim that my account is relevant, not as an identification of marks of the concept's history, but rather as a description of how it actually functions today.

I do not claim that the meanings of knowledge attributions are insensitive to the particular social environments in which they are made. On the contrary, by accepting

[111] Kusch 2011, p. 6.

a form of convention contextualism, I shall include sensitivity to the idiosyncrasies of social structures as a crucial condition in my conceptual synthesis. Therefore, while my account does not *rely* on any one particular claim about how changing social environments have changed the meaning of knowledge attributions, it *implies* that such a process must have occurred. A conceptual synthesis only needs to rely on a genealogy if it turns out that its hypothesis does not and cannot meet the generality requirement, but that rather the needs themselves have changed throughout history. However, since my hypothesis is a generalization of Craig's hypothesis, it has an equally good chance of satisfying the generality requirement. Thus, I argue that my hypothesis relies on nothing but actual evidence of the need for epistemic authority.

It may still turn out that it cannot describe all the ways in which the concept of knowledge is used; that is, it may succeed only in accounting for one very widespread pattern of the use of knowledge attributions. If that is the case, then I have no objections to that limitation and gladly accept Kusch's defense of the undiminished relevance of such analyses. However, considerations of simplicity urge us to consider the stronger claim, that various patterns of use all serve one more general function. While the concept of the subjectively good informant obviously could not play such a broad overarching function, the concept I synthesize would have a much stronger chance. I will not argue for the universality of my final synthesis; if I do manage to properly characterize any use of the concept of knowledge, this synthesis has proved worthwhile.

3. Hypothesis: The need for epistemic authority

In the following chapters, I present my own conceptual synthesis of knowledge. I follow Craig's four step method as outlined in 1.b: hypothesis, explication, synthesis, and comparison. For the reasons given in 2.c, I employ neither a state of nature nor a genealogy.

a. An objectivized hypothesis

Craig's hypothesis fails, because the concept of knowledge does not meet the need for a good informant. Rather, the concept of knowledge meets the need for an *objectively* good informant. Craig mitigated the failure by modifying the synthesized concept. I suggested that we should rather modify the original hypothesis.

However, the need for an objectively good informant might seem to be a far less plausible starting point. It was easy to imagine what would satisfy an inquirer looking for a good informant. How do we explicate what would satisfy an inquirer looking for an objectively good informant? We might have trouble finding intuitively obvious answers, because as we have seen, the inquirer-informant relationship involves society only by hypothetical extension; the basic need is a private one. As I imagine that it will be easier to base the synthesis on a basic need, we should begin with a hypothesis that already acknowledges social circumstances.

As we have seen in section 2.a, people do not only need to be able to find good informants for themselves, they also need to be able to talk about good informants and recommend them to others. The basic goal then is for whole *communities* of inquirers to get true beliefs, now and in the future. We saw above how this goal demands a counterfactual condition; it must be the case that, were the subject to inform an inquirer, he *would* tell the inquirer the truth. In the actual world, this counterfactual assertion condition boils down to an actual true belief condition. Likewise, when an inquirer uses an inquiry stopper, her goal is to stop with a true belief. For her too, we must explicate the true belief condition in order to accommodate her goal scenario.

With the true belief condition, we accommodate the *goal* for both the inquiry stopper and the good informant function. What about the *strategy*? The indicator property takes care of the inquirer's strategy, by ensuring the informant's

trustworthiness. The inquirer looks for informants with indicator properties to reach her goal of finding informants with true beliefs. However, when determining who is trustworthy enough, an individual inquirer asks only whether the inquirer is trustworthy enough for her *own* circumstances. When we begin our hypothesis by considering social and diachronic needs as well, this attitude no longer suffices. People need a concept that can refer to a category of trustworthiness that is relative to their entire community, rather than to themselves as individuals. Thus, I hypothesize that knowledge meets the strategic need for a category of stable trustworthiness.[112]

Craig introduces this stable category as the concept of the objectively good informant. I suggest that we phrase the hypothesis rather by saying that people need to determine which subjects have *epistemic authority*.[113] There are several advantages to this substitution. In the following section, I point out five.

First, we are already accustomed to taking entire societies into account when applying the concept of authority. Thus, the epistemic authority hypothesis can accommodate the objectivization modifications more intuitively than an extension of the good informant hypothesis. As we have seen, an objectively good informant does not actually have to be informing anyone at all. Luigi *would* be able to tell Mario lots of information about the sewage system, but right now Luigi is asleep. It strains our intuitions to claim that Luigi is nonetheless an objectively good informant about the sewage system, because that concept is clearly artificial. On the other hand, it is easy to imagine someone who has authority, despite the fact that he is not exercising that authority. While authority does involve certain kinds of relationships, we more readily understand those relationships as hypothetical. We can say that, in her religion, Barb has the authority to marry couples. Thereby we make a claim about the relationship between Barb and a hypothetical couple. If there was a pair before her hoping to marry, Barb *would* have the authority to marry them. Barb's authority is not affected by the abundance or lack of actual couples, nor by whether she is awake or asleep. Thus, when we consider the need people have for a predicate that involves hypothetical rather than real testimonial relationships, it is more natural to phrase the need as for epistemic authority, rather than for an objectively good informant.

[112] The concept of knowledge combines both a goal need and a strategic need; see section 1.d.iii.

[113] For a discussion of the role of epistemic authority in testimony see Keren 2007. For a connection of Craig's work to epistemic authority see Lane 1999. For a discussion of social and ethical features of epistemic authority and their connection to Craig see Fricker 1998 and Fricker 2007.

The epistemic authority hypothesis also draws attention to a new condition that we must explicate: people need to recognize stable *social groups* of people with good epistemic *statuses*, such as, for example, graduates from an educational program. Since all human societies need some form of education in order to prosper, we may assume that they all need a concept like epistemic authority to label their successful students and make them recognizable to others. A student who studies for and then passes a physics exam is recognized as having attained the authority to answer questions about physics. The same holds for the whole range of academic recognitions, from diplomas to the Nobel Prize. Professions are also markers of stable epistemic statuses. A society recognizes that its hunters are authorities on which weapon should be used to kill a buffalo, while its generals are authorities on which weapon should be used to kill a soldier. Nothing in this status implies any more than what has already been described in explicating the concept of the objectively good informant. Nonetheless, when talking about a good epistemic social status, the concept of the epistemic authority figure is a more intuitive choice than the concept of the objectively good informant.

The explicated demand for epistemic social statuses illuminates another intuitive implication of epistemic authority: its connection to honor. To recognize someone's authority usually involves honoring them. The way in which we recognize a subject's *epistemic* authority is, on this hypothesis, by recognizing the subject as someone who can answer a certain question.[114] Once again, this condition is already implicit in the good informant hypothesis. The epistemic authority framework just gives us a handy tool for imagining how such honoring works; we are used to connecting authority to honor. Furthermore, this explicated condition should result in a closer match to knowledge, for intuitions that being recognized as a knower involves honor are widely spread.[115] Of course, not every knowledge attribution implies honor, but then neither does every recognition of authority. Authorities on underwater basket-weaving have a notoriously tough time garnering respect. We may also expect that certain kinds of authority, such as the epistemic authority to say what time it is, are too easy to come by, and accordingly ill-deserving of honor. Nonetheless, the level

[114] Thus we need not try to find some particular *faculty* that deserves honor, as Lackey assumes in Lackey 2007. Of course, Lackey would worry about *merely* lucky people, who are only accidentally trustworthy and thus deserve no honor. Such an objection will be addressed in section 6.c.
[115] Kusch 2009, pp. 81 – 82.

of honor involved might still be *negatively* visible; people who are unable to read a clock might well have a recognizably low epistemic status in their community, and experience shame as a result. In any case, it seems true for both authority and knowledge that, the higher the investment, the greater the honor awarded.

Furthermore, thinking of knowledge as a response to a need for epistemic authority helps to highlight the *normative* aspects of knowledge. As Melissa Lane points out,

> I am rationally bound to accept certain persons as knowers on the basis of their socially acknowledged standing as such, and I can be held socially and epistemically liable if I do not.[116]

All social systems are maintained by a network of expectations and demands; there is no reason for the system of epistemic evaluation to form an exception. Therefore, the concept we are synthesizing must include conditions about who *ought* to be believed and who ought not to be believed. Subjects with epistemic authority ought to be believed. An inquirer who does not believe a subject with epistemic authority is violating an epistemic norm. If a sexist inquirer who only trusts women meets a male math whiz, then she would be epistemically blameworthy not to believe his math-related assertions.[117] He has the requisite epistemic authority; therefore she ought to believe him. Having authority obviously entails normative demands on others; it is not intuitively as clear that being an objectively good informant would entail the same. Thus, the normative condition demanded by the modified hypothesis also favors speaking of epistemic authority rather than an objectively good informant.

In the following sections, I continue to elaborate on some facets of the need for epistemic authority, particularly as it relates to practical reasoning. However, I do not stipulate a precise definition of the term, preferring to leave it vague and suggestive. To clarify it, we would have to turn to the concept of authority in general, of which epistemic authority is presumably one particular kind. The more we learn about the need for a concept of authority in general and its function in society, the more we may expect to learn about the need for epistemic authority in particular and the way in which the concept of knowledge meets that need. As I am not in the position to

[116] Lane 1999, p. 220.

[117] This would be a case of epistemic injustice, a phenomenon thoroughly explored by Miranda Fricker. (Fricker 2007) I have said the sexist inquirer would be epistemically blameworthy; Fricker argues that epistemic injustice is both an epistemic and an ethical vice.

offer a full synthesis of the need for authority in general, there are bound to be aspects of our need for epistemic authority that I overlook. However, I take these inevitable oversights as a strength, rather than a weakness, for they illuminate potentially fertile avenues of further research.

In sum, when we *begin* with the hypothesis that knowledge meets the need for a stable category of trustworthiness, we may more naturally speak of the need for epistemic authority rather than the need for an objectively good informant. I have suggested that the alternative hypothesis offers five advantages: it is more plausibly construed as relative to hypothetical relationships, it implies a social status, it implies honor, it implies a normative aspect, and it suggests further questions by connecting knowledge to the broader concept of authority.

b. Continuity with the subjective hypotheses

In the section above, I identify five advantages for the alternative hypothesis. However, I also claim that my hypothesis *subsumes* Craig's own original hypothesis, as well as Kelp and Kappel's. The needs upon which their syntheses rest are specific instances of the more general need for a category of stable trustworthiness.

An inquirer may need a subject with epistemic authority, simply because subjects with epistemic authority are trustworthy and an inquirer needs to find out whom she can trust when searching for good information. In other words, people need subjects with epistemic authority in part because they need good informants. A good informant is a subject with epistemic authority, and the need for a concept of the good informant is one instance of the more general need for a concept of epistemic authority. If I am busy picking berries in a tiger-ridden wood, I will not take any madman's ravings as a reason to flee; I will only trust the warning of someone who seems to possess the requisite authority on the question whether or not a tiger is coming. Being up a tree is a fine way to establish just that authority. To that extent, 'has legitimate epistemic authority on the question whether *p*' is simply another way of saying 'is a good informant as to whether *p*'.

Epistemic authority also covers the inquiry stopper function hypothesized by Kelp and Kappel. Inquirers may only cease inquiry as to whether *p* if they have reached the point where they have the requisite authority to believe that *p*. If I want to find out whether a ladder is sturdy, I will perform whatever tests are necessary to

make myself an authority on the matter of that ladder's sturdiness. Again, all this only goes to say that 'has legitimate epistemic authority on the question whether *p*' may express the same idea as 'may permissibly stop inquiry as to whether *p*.'

I assume that no one's intuitions have already risen in rebellion. A particularly clear example of 'epistemic authority' being used in the two senses outlined above may be found in the 2005 paper *Say's Who?: Epistemic Authority Effects in Social Judgment*,[118] which studies the make-up of epistemic authority with the tools of experimental psychology. The authors introduce their topic as follows:

> As we negotiate our way through the labyrinths of interpersonal relations and task exigencies, we encounter a continuous flow of information in the form of communications, advice, exhortations, and pleas from a variety of sources. These pose the ubiquitous question of whom to (informationally) trust and whose statements to discount or regard with suspicion.[119]

They then present our assessment of others' epistemic authority as necessitated by this state of affairs. As such, they make 'having epistemic authority' analogous to 'being a good informant.' They go on to define epistemic authority more closely by using the considerations that motivate Kelp and Kappel:

> In its role of a ''stopping mechanism,'' a source's authority plays, therefore, a motivational role in information processing determining the amount of energy the individual is prepared to devote to continued epistemic activity in a domain.[120]

Thus, they equally understand 'having epistemic authority' as analogous to 'being able to stop inquiry'.

Furthermore, they specifically stress the possibility that an inquirer may serve as her *own* epistemic authority.[121] If a subject doesn't feel confident about her own expertise in a certain subject, then she will tend to seek advice from other people before making a decision as to whether *p*. In such situations, the inquirer looks for the assertion of a trustworthy informant before committing herself to belief. She may not believe that *p* on her own authority. If on the other hand, she believes that she herself "knows best,"[122] then she will simply take her own belief that *p* as

[118] Kruglanski, et al. 2005.
[119] Ibid., p. 346.
[120] Ibid., p. 352.
[121] Ibid., pp. 370 – 376.
[122] The authors found the formulation "Who knows best?" to be the question that most clearly evoked the notion of epistemic authority. (Ibid., p. 357)

authoritative enough. Thus epistemic authority can function in the first-person as well as in the third-person.

This example from non-philosophical literature should suggest that equivocating between the concept 'epistemic authority' on the one hand and the concepts 'good informant' and 'inquiry stopper' on the other hand is no rash stipulation, but is rather supported by generally shared intuitions about what authority in the cognitive domain entails. The idea of epistemic authority can capture our intuitions about both the good informant function and the permissible inquiry stopper function.

c. A problem anticipated

I have introduced a new hypothesis: the concept of knowledge meets the strategic need for a concept of epistemic authority. I claim that this hypothesis takes all the concerns of objectivization into account from the get-go. However I have further claimed that the hypothesis covers the function of the subjectively good informant and the subjectively permissible inquiry stopper as well. The result is appealing in its comprehensiveness, but is it tenable?

The concept of the good informant serves a vital function thanks to its sensitivity to inquirers' personal needs. The concept of stable trustworthiness serves a vital function thanks to its insensitivity to same. How can the concept of epistemic authority include both functions? This contradiction threatens to prevent a coherent synthesis. The conditions explicated as necessary to one function may turn out anathema to the other. To support my claim that one single concept can play both functions, I must explicate a condition which allows for both enough sensitivity and enough insensitivity.

Just such a condition is in fact available, as I show at the end of Chapter Five. Before reaching this point though, I must elaborate on the kind of sensitivity we should explicate from the hypothesis. This will be the theme of Chapter Four.

4. Explication: Sensitivity to practical context

a. Knowledge characterized by its authoritative role in practical reasoning (KAP)

i. The vague reliability condition

In the last chapter, I introduced a new hypothesis: knowledge meets the strategic need for epistemic authority. This hypothesis should accommodate the demands of objectivization. I have already explicated three new conditions from the hypothesis: a social status condition, an honor condition, and a normativity condition. I will return to the importance of these conditions in section d of this chapter. For the moment, however, I set them aside in order to concentrate exclusively on the function of the reliability condition. With the provisional exclusion of the new social conditions, we may simply make a couple substitutions to the objectivized concept of the good informant,[123] and explicate this synthesis:

Knowledge as true and potentially Authoritative belief (KA)
S_1 knows relative to S_1 - S_n iff *p*, S_1 believes that *p*, and S_1 has a property X which correlates with being right about *p* reliably enough that S_1 - S_n may appropriately use the belief that *p* on S_1's authority.

Before we go on to step four and compare KA to the concept of knowledge, we must determine just how reliably the property X must correlate with being right about *p* – how reliably is 'reliably enough'? What standards of reliability will meet people's need for epistemic authority? Once we explicate the relevant standards, we will be in the position to clarify what makes it 'appropriate' for an inquirer to believe on an informant's authority.

ii. Two sources of inspiration

To explicate the reliability condition, I will borrow an idea from John Hawthorne and Jason Stanley's subject-sensitive invariantism, namely, that knowledge is the norm

[123] See 2.a.vi

of practical reasoning.[124] According to Hawthorne and Stanley, one should only use *p* as a reason for acting if one knows that *p*. Inspired by this account, I will argue that S_1 may appropriately have epistemic authority over $S_{1\text{-}n}$ regarding *p* only if it would be acceptable for $S_{1\text{-}n}$ to use the belief acquired on S_1's authority as a reason for acting.

Due to the testimonial needs included in my hypothesis, however, my position will not mimic subject-sensitive invariantism. To develop it properly, I must rather borrow an idea from an opposing camp: contextualism. Namely, I make use of the contextualist idea that the meaning of a knowledge attribution is relative to the person making the attribution. I go through the differences between my account and Hawthorne and Stanley's account in more detail below, in section 4.c.

In sum, I will argue that the meaning expressed by a knowledge attribution is directly affected by the practical environment of the attributor. Thus, I use an element of subject-sensitive invariantism in order to motivate a contextualist account of knowledge. Most contextualist accounts of knowledge have been motivated by the desire to properly interpret skeptical arguments, explaining why they are both compelling and unthreatening to everyday knowledge claims.[125] My contextualist synthesis of knowledge, on the other hand, is focused entirely on the functional role the concept of knowledge plays in social life, and not at all with skepticism.

iii. Craig's prototypical case

Craig describes the concept of the good informant as a response to the demands of practical reasoning. In order to get all the information she needs to successfully pursue her practical goals, the inquirer needs to detect good informants. Once she finds a good informant, she asks him whether *p* or not *p*. If he asserts that *p*, she can believe *p* on his authority and accordingly use *p* as a premise in her practical reasoning. So, Craig's portrayal of the prototypical case looks as follows, in which S_2 is an inquirer and S_1 an informant:

The prototypical case of the good informant and the inquirer

S_2 needs to find out whether *p*. S_1 truly asserts that *p*. S_2 detects that S_1 has a property

[124] Hawthorne & Stanley 2008.
[125] Willaschek 2007, p. 254.

X which correlates with being right about *p* reliably enough that S_2 believes that *p* on S_1's authority and accordingly uses *p* as a premise in her practical reasoning.

Two points require clarification. First, I must clarify what I mean by using *p* as a premise in practical reasoning. To use *p* as a premise in one's practical reasoning means to act as if *p*. If *p* makes action A preferable, and not *p* makes action B preferable, then a subject acts as if *p* when she performs action A. For example, if it will probably rain, Marion would prefer to bring an umbrella. If it will probably be dry, Marion would rather do without the hassle of lugging an umbrella around; she has many other things to carry and a long way to travel. A weatherman tells her that it will probably be dry. She detects that being a weatherman correlates fairly reliably with being right about such things and accordingly decides to risk the chance of rain, leaving the umbrella at home. Because Marion believes on the weatherman's authority, she uses the proposition "It will probably be dry" as a premise in her practical reasoning. By leaving the umbrella at home, she acts as if it will probably be dry.[126]

Second, I must clarify the relationship between the inquirer's trust and the inquirer's reasoning. In Craig's prototypical case, if an informant is good, then an inquirer may believe *p* on his authority and *accordingly* use *p* as a premise in her practical reasoning. The 'accordingly' should indicate that she uses the belief in her practical reasoning *because* of S_1's authority. This is a causal relationship that I tentatively characterize along David Lewis' account of a causal chain of counterfactual causal relations:[127] If the inquirer had not detected the informant's indicator property, then the inquirer would not have acquired a belief that *p* on his authority. If the inquirer had not acquired a belief that *p* on his authority, then the inquirer would not have used *p* as a premise in her practical reasoning. For example, Fred climbs a tree and starts screaming, "A tiger is coming!" and the inquirer on the ground runs away. If the inquirer had seen Fred eating hallucinogenic mushrooms, then the inquirer would not have believed that a tiger was coming on his authority. If

[126] For further explanation and a host of clarifications, see Fantl & McGrath 2002.
[127] For an introduction to the theory, see Menzies 2009. I assume that the problems with Lewis' theory do not substantially affect my explication. The following discussion should demonstrate that my explication would particularly welcome modifications emphasizing the contextual nature of causality.

the inquirer had not believed that a tiger was coming on his authority, then she would not have acted as if a tiger was coming, i.e., ran away.

A subject is only in truth a good informant if his authority is *appropriate*. A subject's authority is appropriate if and only if an inquirer would be *rational* to use a belief based on his authority as a premise in her practical reasoning. A causal chain goes from the informant's authority over whether *p*, to the inquirer's belief that *p*, to the inquirer's action as if *p*. In order to evaluate the informant's authority we must ask whether the first link makes the final link rational. If the level of the reliability of the informant's indicator property makes the inquirer rational to act as if *p* based on the inquirer's trust in the informant, then the informant's authority is appropriate. If on the other hand, the inquirer would be irrational to use a belief acquired out of trust in the informant as a premise in her practical reasoning, then the informant's authority is not appropriate. Imagine Luke tells Colleen that a tree branch is strong; Colleen trusts him and crawls out onto the branch. If her crawling out onto the branch is rational, then Luke has the epistemic authority to make the assertion. That is, if Colleen may use "the tree branch is strong" as a premise in rational practical reasoning based on her trust in Luke, then Luke's belief that "the tree branch is strong" really is authoritative.

Thus, the appropriateness of S_1's authority for S_2 that *p* is derivative from the rationality of using *p* as a premise in S_2's practical reasoning. The amount of reliability required for using *p* as a premise in rational practical reasoning dictates the amount of reliability S_1 requires to have authority as to whether *p*. We may describe the concept of the good informant with the following biconditional:

The concept of the good informant characterized by its authoritative role in practical reasoning

S_2's belief that "S_1 is a good informant that *p*" is true iff *p*, S_1 asserts that *p*, and S_2 detects that S_1 has a property X which correlates with being right about *p* reliably enough that S_2 believes that *p* on S_1's authority and accordingly uses *p* as a premise in rational practical reasoning.

iv. Objectivizing everything besides the reliability condition

The previous section clarifies the relationship between being a good informant and practical reasoning. However, I have rejected Craig's good informant hypothesis, arguing that we must begin with the epistemic authority hypothesis instead. Nonetheless, I have also argued that the epistemic authority hypothesis can subsume the good informant function, doing all of its work in addition to its further social functions.

I will argue that the concept of knowledge can do the work of the good informant concept by playing the role in practical reasoning just outlined. At the end of the chapter, I will acknowledge the tension between this subjective feature and the social conditions explicated above.[128]

Losing the subjective conditions (1) and (4) allows that S_1 need not actually tell S_2 that *p* in order to count as knowing that *p*; as we saw above, the social hypothesis demands instead that S_1 believe that *p*. Losing the subjective condition (2) allows that S_2 need not actually detect S_1's indicator property. In addition, I take Kappel's inquiry stopper hypothesis into account, and his maneuver of combining the two functions by substituting S_2 with $S_1 - S_n$. We end up with the following synthesis:

Knowledge characterized by its Authoritative role in Practical reasoning (KAP)
S_1 knows that *p* relative to $S_1 - S_n$ iff *p*, S_1 believes that *p*, and S_1 has a property X which correlates with being right about *p* reliably enough that $S_1 - S_n$ could believe that *p* on S_1's authority and accordingly use *p* as a premise in rational practical reasoning.

According to KAP, the concept of knowledge functions to label subjects who are in such a good epistemic position that someone could rationally rely on their beliefs when deciding what to do. Refusing to attribute knowledge to a subject functions as a sign that someone would have to find better evidence before acting as his beliefs suggest. So, if someone would be rational to believe the branch was strong on Luke's authority and accordingly crawl out on it, then Luke knows that the branch is strong. If someone would be irrational to do so, then Luke does not know that the branch is

[128] As anticipated in 3.c.

strong. We can use this hypothetical relationship to judge whether or not Luke knows that the branch is strong even if Luke is alone and no one wants to climb the tree.

v. KAP vs. Hawthorne: a foreshadow

By connecting the appropriateness of S_1's authority to the rationality of S_1 - S_n's practical reasoning, I clarified how reliable a subject must be to count as knowing. This tactic is inspired not only by Craig, but also by the work of John Hawthorne. In his book *Knowledge and Lotteries*, he writes,

> Insofar as it is unacceptable – and not merely because the content of the belief is irrelevant to the issues at hand – to use a belief that *p* as a premise in practical reasoning on a certain occasion, the belief is not a piece of knowledge at that time.[129]

Hawthorne's formulation is simpler than KAP. If we were to reformulate his principle in terms of causality, his causal chain would only have one link:

> If S did not believe that *p*, then S would not use *p* as a premise in her practical reasoning.

KAP explicitly includes an additional link to begin the chain:

> In the first person case, if S did not have a property that indicated authority as to whether *p*, then S would not believe that *p*.
> If S did not believe that *p*, then S would not use *p* as a premise in her practical reasoning.

Of course, Hawthorne would probably accept that that the concept of knowledge plays a function in that first link, but only in the first-person form. KAP, on the other hand, also allows for the following chain:

> In the third person case, if S_1 did not have a property that indicated authority as to whether *p*, then S_2 would not believe that *p*.
> If S_2 did not believe that *p*, then S_2 would not have used *p* as a premise in her practical reasoning.

In this chain, Hawthorne could only recognize a function for the concept of knowledge in the second link: it serves to separate beliefs one should act on from beliefs one should not act on. He would be forced to deny that knowledge can function in the first link as a way to separate testimony one should believe (and accordingly act on) from testimony one should not believe (and accordingly act on).

[129] Hawthorne 2004, p. 176.

Therein lays the greater explanatory power of KAP as opposed to Hawthorne's version. I return to this comparison at much greater length in sections 4.c and 6.c.

vi. Crisp's bizarre counterexample

Thomas Crisp constructs a thought experiment which he claims presents a counterexample to such a conception of knowledge.[130] In his thought experiment, Dr. Evil is monitoring your thoughts and has threatened to torture and kill you if you use '2+2=4' as a premise in your practical reasoning.[131] In that case, it would not be practically rational to do so, yet surely you still know that '2+2=4'. Thus, we have a scenario in which a piece of knowledge may not be rationally deployed in practical reasoning. Therefore, Crisp reasons, the reliability condition presented above cannot be necessary to knowledge.

To meet this objection, we must return to the considerations of section 1.d.i. If we may conceive of the concept of knowledge by way of explicating conditions from generally shared needs, we must include all and only those conditions which *usually* meet those needs. Bizarre possibilities do not suffice to affect the synthesis. In nearly all cases, a true belief acquired on appropriate authority is rational to use in practical reasoning. Crisp needs to demonstrate that there is a set of ordinary cases in which that doesn't hold. His bizarre example is not enough to demonstrate that the concept of knowledge is unaffected by its role in practical reasoning.

b. Explicating contextual attributor sensitivity

i. Practical environments change

People need to recognize subjects whose epistemic position is such that *someone* would be able to act on the authority of their beliefs. So if, in a nearby possible world, Colleen would be rational to act on Luke's belief that the tree branch is strong, then Luke counts as knowing it, even if Luke is actually alone and no one wants to climb the tree. What happens, though, in situations like the peanut example?[132] Nelville believes that his cookies have no peanuts in them, and he has

[130] Crisp 2005.
[131] Ibid., p. 139. Crisp assures us that, although Dr. Evil is monitoring your thoughts by means of an implanted chip, your cognitive abilities have not been impaired.
[132] See section 2.c.i.

fairly good, though not absolutely decisive, evidence for the belief. Estelle would be rational to act on his belief by eating a cookie, because she would only risk experiencing a taste she dislikes. Ida would not be rational to act on his belief by giving one of the cookies to her child, who is fatally allergic to peanuts. So, does Nelville know or not? According to KAP, Nelville would know relative to Estelle, but would not know relative to Ida. KAP cannot provide one single answer, because whether a premise may be used in rational practical reasoning or not varies from reasoner to reasoner. Specifically, it varies relative to a reasoner's *practical environment*. The notion of a practical environment is taken from Hawthorne.[133] It consists of any factors that affect which premises one may use in rational practical reasoning.

Some factors that affect a practical environment are already traditionally accounted for in epistemological theories. The amount of evidence one has for a belief, for example, can affect whether or not it is rational to act on it. In what follows, I am more interested in the factors which a practical environment includes, but which have usually been *excluded* from analyses of knowledge. In the peanut example, Nelville's amount of evidence stays constant; what changes is the risk involved in acting on his belief. By explicating the concept of knowledge as that which plays the role in practical reasoning described by KAP, I assume that such blatantly practical considerations can affect whether or not a subject knows.

The rest of this section introduces three such factors that affect a practical environment. Furthermore, in presenting examples of these factors, I also demonstrate how knowledge can function by virtue of its sensitivity to the relevant practical factor. In so doing, I already anticipate the fourth step of comparison, eliciting intuitions that it is indeed acceptable for knowledge to be applied just as my explication deems it would be.

(1) End of inquiry

Rational actions are directed towards some sort of end. By changing the end of inquiry, the sorts of premises that one may use in rational practical reasoning also change.

[133] Hawthorne 2004, pp. 173 – 180.

Example: The book and the ticket[134]

Buying a Guide Book

S_2 is in a book store, looking to buy a guide book. She can only afford one and must choose between a guide book for a distant land and a local guide book.

END OF INQUIRY: Decide which book to buy.

S_1 tells S_2 "You won't be able to afford to travel to that distant land."

S_1 is S_2's financial advisor. S_2 believes that S_1 knows that she won't be able to afford to travel to that distant land. As a result, S_2 makes the rational decision to buy the local guide book.

Selling a Lottery Ticket

S_2 has bought a ticket in a big lottery for a dollar. Someone offers to buy the ticket off her for a penny.

END OF INQUIRY: Decide whether or not to sell the ticket.

S_1 tells S_2 "Your ticket will lose."

S_1 doesn't have any inside information about this particular lottery. S_2 believes that S_1 does not know that the ticket will lose. As a result, S_2 makes the rational decision to keep her ticket.

Changing the end of inquiry changed whether it was acceptable for S_2 to use the premise "Your ticket will lose" as a premise in rational practical reasoning. The end of inquiry is a factor in a practical environment.

Furthermore, the example reveals interesting intuitions about knowledge attributions. Imagine that the lottery's prize is big enough that, were S_2's ticket to win, she would in fact be able to afford to travel to that distant land. In that case, the truth of S_1's assertion in Buying a Guide Book would entail the truth of S_1's assertion in Selling a Lottery Ticket. Nonetheless, it's generally agreed that S_2 is in fact acting rationally and in her own best interest when she accepts S_1's testimony in Buying a Guide Book and rejects it in Selling a Lottery Ticket.[135] According to my explication,

[134] This example is adapted from Hawthorne 2004.

[135] Janet Levin argues for a different interpretation of the lottery case, arguing that since buying a lottery ticket in the first place is so irrational, it is not rationality that bars one from saying that the

the concept of knowledge would thus apply to S_1 in Buying a Guide Book, but not in Selling a Lottery Ticket. That is, the end of inquiry would have to be able to affect what S_1 knows about S_2's financial situation. Since the truth of S_2's knowledge attributions in these examples is intuitively acceptable, the contextualist sensitivity I have explicated promises to do well in the fourth step of comparison.

(2) Stakes

As already shown in the peanut example above, raising the stakes of being wrong about a certain premise can change whether or not it is acceptable to use that premise in rational practical reasoning.

Example: "John was at work."[136]

Thelma and Louise made a small office bet as to whether their often absent colleague John would be at work that day. Before leaving the office, Louise sees John's hat hanging on the coatrack and hears Fred call out "Why don't you clear that letter with John quick before you send it off?"

Petty Bet

At a bar, Louise runs into Thelma. Thelma had bet five dollars that John would be absent. If Thelma believes that John was at work when he wasn't, she loses five dollars that she otherwise would have kept. If she loses the five dollars, she won't be able to afford another drink, which she mildly desires.

STAKES: Low

Louise tells Thelma that John was at work, describing what she saw and heard.

Thelma believes Louise and as a result, makes the rational decision to give Louise five dollars.

Criminal investigation

Later that night, Louise is visited by the police, who are investigating a terrible crime. If the police believe that John was at work when he wasn't, they might arrest an innocent man, thereby endangering lives and their jobs.

---

ticket will lose. Rather the irrational reason for buying the ticket also prevents the ticket holder from coherently believing that it will lose. (Levin 2008, pp. 372 – 373)

[136] This example is adapted from DeRose 2009.

STAKES: High

Louise tells the police that John was at work, describing what she saw and heard. The police believe Louise and as a result, make the rational decision to continue looking for a reliable witness as to whether John was at work.

Louise's position in relation to her assertion that "John was at work" doesn't change. However, it is only acceptable as a premise in rational practical reasoning in a fairly low stakes situation. The height of stakes is a factor in a practical environment.

(3) Relevant alternatives

To be an authoritative informant that *p*, a subject must be able to rule out certain possibilities which entail that not *p*. These possibilities, or 'counterpossibilities,'[137] may be called the relevant alternatives. Fred Dretske first introduced this feature of knowledge claims and defined it as follows,

> A relevant alternative is an alternative that might have been realized in the existing circumstances if the actual state of affairs had not materialized.[138]

S_1 cannot responsibly believe that S_2 knows that *p* unless S_1 believes that S_2 has ruled out a certain set of relevant alternatives that would entail not-*p*.[139]

Example: "Lefty killed Otto"[140]

Did Lefty do it?

Altin tells Deia that "Lefty killed Otto." Deia is doubtful; she asks, "Are you sure it wasn't Righty who killed Otto?"

---

[137] Bach 2005.

[138] Dretske 1970, p. 1021.

[139] One way to integrate relevant alternatives into an account of knowledge is to consider knowledge attributions as based on contrastive questions: "Does S know that *p*, or is S unable to eliminate alternatives 1, 2, and 3?" See Blaauw 2005 for the thesis that knowledge relations hold between three relata: a subject S, a proposition p, and a set of contrastive propositions Q. While he presents the view as an alternative to contextualism, that is due to his identification of contextualism with the claim that 'knowledge' is an indexical. I conceive of contextualism more loosely as any account which claims that knowledge is relative to the attributor's context and so would construe Blaauw's account as another form of contextualism.

[140] This example is adapted from Dretske 1970. The more commonly cited example from the same paper concerns zebras in a zoo: must one be able to rule out that it is a cleverly disguised mule in order to know that it is a zebra? If so, skeptical implications abound. Since my interest is not in skepticism, I have chosen a somewhat more quotidian example.

Relevant Alternative: Righty killed Otto.

Altin answers, "Yes, I was playing cards with Righty when it happened. He definitely didn't do it." Accordingly, if Deia had to use the premise that Lefty killed Otto in her practical reasoning, she could rationally do so on Altin's authority.

Is Otto dead?

Altin tells Deia that "Lefty killed Otto." Deia is doubtful; she asks, "Are you sure Otto really died?"

Relevant Alternative: Lefty only injured Otto.

Altin answers, "Yes, I spoke to the coroner myself." Accordingly, if Deia had to use the premise that Lefty killed Otto in her practical reasoning, she could rationally do so on Altin's authority.

If Altin is to count as knowing that "Lefty killed Otto," he must be able to rule out the relevant alternatives. Which alternatives are relevant depends on the context of inquiry; in the first example Deia demands that he be able to rule out a different culprit, in the second, Deia demands that he be able to rule out a mere injury. Why should Deia's demand change like that? There are several possible answers. We can already see from the previous two examples that both the end of inquiry and the stakes of a practical environment can affect which alternatives are relevant.[141] In "Selling a Lottery Ticket," the very slim chance that S_2's lottery ticket will win is a relevant alternative. S_2 can't accept S_1's assertion as a reason for acting unless she can rule out that very slim chance. In "Buying a Guide Book" the very slim chance that S_2's lottery ticket will win is not a relevant alternative. S_2 can accept S_1's assertion as a reason for acting, despite the fact that she hasn't ruled it out. Likewise, in "Petty Bet" the possibility that John left his hat there the day before and that Fred just made a mistake is not a relevant alternative, but in "Criminal Investigation" its feasibility prevents the police from using Louise's testimony as a reason for prosecution.

[141] Jonathan Schaffer argues that all the other factors may be reduced to the factor of relevant alternatives. (Schaffer 2005)

What else can make an alternative relevant? Is it enough that an alternative is salient to the speaker?[142] Or are there more 'objective' criteria for relevance? My answer to this question must await a clearer model of practical environments, which I offer in section 5.e.

The three factors outlined here leave many crucial details unaddressed. The one general point which should already be clear is that, if the concept of knowledge is to play the role described by KAP, it must be sensitive to the end of inquiry, the stakes, and the relevant alternatives in the potential inquirer's practical environment.

ii. "Knowledge" changes according to context

As the examples given above demonstrate, when the factors of a practical environment change, then the acceptability of *p* as a reason to act changes. Given KAP, the acceptability of *p* as a reason to act directly affects whether knowledge may be truthfully attributed. The subject counts as knowing if *someone* could rationally act on his belief. The subject counts as not knowing if *someone* could not rationally act on his belief. Yet both circumstances can be simultaneously true. Estelle could rationally act on Nelville's belief; but Ida could not. Whether or not "Nelville knows that the cookies have no peanuts" is a true statement depends on whether the speaker who states it is considering Nelville's potential relationship with Estelle or his potential relationship with Ida.

In other words, KAP supports *contextualism* about knowledge. According to the contextualist portrayal of knowledge, there need not be one single right answer to the question "Does Nelville know?" Estelle can truthfully say "Nelville knows that *p*" while Ida can truthfully say "Nelville doesn't know that *p*" in reference to the same Nelville and the same proposition at the same time. The two sentences do not necessarily contradict one another because the meaning of the word 'know' and thereby the actual proposition expressed can *change*. That is, it is possible for Estelle to mean "Nelville can rule out relevant alternatives 1, 2, and 3," while Ida means "Nelville can rule out relevant alternatives 1, 2, 3, 4, and 5." Since they both mean different things by the word 'know', both sentences can be true without entailing a contradiction. Just as S can be simultaneously both 'here' and 'not here' depending

[142] I consider an account that answers yes to this question in section 5.c.

on who is talking, S can simultaneously both 'know' and 'not know' that *p* because the meaning of the word 'know' can shift according to who is talking.

The context sensitivity of "here" is of course much more obvious than the alleged sensitivity of "know(s)". Contextualists traditionally compare "know(s)" to more subtly context sensitive words like "tall". One person can truthfully say "S is tall" while another can truthfully say "S is not tall" in reference to the same person at the same time. The two sentences need not contradict, because the semantic value of the word "tall" can be different in two different attributors' context. That is, it is possible that the first person means, "tall for a fourth grader" while the second means "tall for a basketball player."

The concept of knowledge cannot function to meet the hypothesized need unless it is sensitive to an implied practical environment. As a result, I explicated a contextually sensitive reliability condition, as described in the formula KAP. According to that condition, the meaning of the word "knowledge" can change according to the speakers' context. Thus, if KAP is acceptable, the concept of knowledge needs just the kind of flexibility that contextualists claim it has.

iii. Functional contextualism

Contextualism about knowledge is generally taken to be a thesis about the semantic meaning of the word 'know'. One can thus be an invariantist about the semantic meaning of the word 'know,' while still claiming that the pragmatic implications of the word 'know' vary with context. Indeed, Unger, who first introduced the terms 'invariantist' and 'contextualist', assumed throughout that all epistemologists must agree to some sort of contextualist shift in communicated meaning.[143] He portrayed the vital difference as whether the shift happens in the semantics or the pragmatics of the knowledge attribution. That is, does the *proposition* expressed by a knowledge attribution, and thereby its truth conditions, change according to context? Or is it rather what is *implied* and understood beyond the propositional content that changes according to context?

For example, if I say, "Nobody will be laid off," the semantic meaning of my statement is just that: "Nobody will be laid off." If someone halfway across the world is laid off, my statement expresses a semantic falsehood. However, it might be very

[143] Unger 1984.

well understood that I do not actually mean "Not one single person will be laid off" but rather "Nobody *in this company* will be laid off." The limited reference class is not part of the semantics of my sentence, but rather a pragmatic implication. So what I communicate and what my interlocutor understands is not the actual proposition I *said*, but rather the one I *meant*. For the concept of knowledge we can understand the difference as follows: either the proposition expressed by 'knowledge' changes according to practical environment, or what people mean when they use the word 'knowledge' changes according to practical environment. If the latter is the case, then people who, for example, say "S knows that *p*" in a low stakes environment might be saying something false, but meaning something true, such as "S is very close to knowing that *p*."[144]

What does a functional synthesis have to say about the distinction between pragmatics and semantics? As Craig recognizes, not much.[145] A semantic invariantist who allows for contextual pragmatic implications may still argue that the *use* of the word 'knowledge' fills the function described by KAP. Of course, she would be left with the burden of demonstrating why the semantics of the word diverge from the function of the word, but such a demonstration might not be impossible. While a functional synthesis must compare the explicated concept to the empirical use of the word, it is not limited to semantic data. If the word can fill its function thanks to pragmatic implication, the thesis is still supported. Thus the contextualism defended here is functional contextualism, which may be either semantic or pragmatic.

However, I still have to defend my functional analysis against those who hold that, according to empirical data, what is primarily communicated and understood by knowledge attributions varies neither semantically nor thanks to pragmatic implication.[146] Since that is the distinction that interests me here, I shall restrict the label 'invariantism' to those who hold that position. I shall refer to those who argue that the uptake of knowledge attributions can vary thanks to pragmatic implication as *pragmatic contextualists*.

The debate over whether shifts in the meaning of knowledge attributions are semantic or pragmatic involves deciding whether intuitive knowledge attributions are

[144] I look at such an account more closely in 5.d.
[145] Craig 1990, p. 167.
[146] Such as Kent Bach in Bach 2004 and Bach 2005. I discuss his position further in section 5.b.

based on conversational propriety or on truth conditions.[147] By disregarding this debate altogether, I free my hands for a different project. Rather than focusing on particular uses of the word 'know' that could shed light on that distinction, I attempt to achieve a more thorough understanding of the various factors that affect what is communicated through knowledge attributions, regardless of how that communication takes place. The results of this more general account can subsequently contribute to the semantics/pragmatics debate, for once we have a better idea of what kinds of shifts to look for, we should be better able to choose which sentences to linguistically analyze.

c. Not explicating invariant subject sensitivity

In elaborating on how the concept of knowledge might function to meet the need for epistemic authority, I have drawn from Hawthorne and Stanley's exemplary work on the link between knowledge and practical reasoning. However, as foreshadowed, I made several crucial changes. Here, I present the outlines of Hawthorne and Stanley's account of knowledge and explain the three main points on which it differs from KAP. I also describe how the two accounts should be compared. The actual comparison and its verdict must wait until KAP has been more satisfyingly developed, in section 6.c.

i. The Reason-Knowledge Principle

Hawthorne and Stanley claim that knowledge functions as the norm of practical reasoning.[148] They offer the following formulation:

> **The Reason-Knowledge Principle**
> Where one's choice is p-dependent, it is appropriate to treat the proposition that p as a reason for acting iff you know that p.[149]

[147] Brown 2005, p. 266. Rysiew 2007, p. 627.
[148] Hawthorne 2004. Hawthorne & Stanley 2008.
[149] Hawthorne & Stanley 2008, p. 9.

For a choice to be "p-dependent" simply means that the question whether *p* is relevant to the action at hand. Or, in Hawthorne and Stanley's more technical explanation:

> Let us say that a choice between options x1…xn is *p dependent* iff the most preferable of x1….xn conditional on the proposition that p is not the same as the most preferable of x1…xn conditional on the proposition that not-p.[150]

Like KAP, the Reason-Knowledge Principle entails that a change in practical environment will result in a change in the applicability of the concept of knowledge, even if the subject's belief and all the traditional factors, such as how much evidence she has, remain the same. Take "John was at work." In Thelma's practical environment it is appropriate for Thelma to treat the proposition that "John was at work" as a reason for action. Therefore, according to the Reason-Knowledge Principle, Thelma knows that *p*. In the police's practical environment it is not appropriate for the police to treat the proposition that "John was at work" as a reason for action. Therefore, the police do not know that *p*. Despite having identical evidence, the difference in their practical environments entails that Thelma but not the police may know that *p*. This interpretation of the example is similar to the one KAP provides, but it is not identical.

The Reason-Knowledge Principle has three major differences from KAP, all of which are interrelated. First, it is motivated by the evaluation of behavior rather than the evaluation of a belief or assertion. Second, it refers exclusively to first person scenarios, resulting in subject-sensitive invariantism rather than contextualism. Third, it refers to knowledge per se rather than knowledge attributions.

ii. Contrasting hypotheses

The Reason-Knowledge Principle is motivated by behavior appraisal scenarios. If we were to rephrase their work in the terms of a conceptual synthesis, we could say that their principle is based on the hypothesis that the concept of knowledge functions to meet the need to appraise behavior. People need the concept of knowledge so that they may exercise epistemic judgment over the behavior of others. That is, on the Hawthorne and Stanley account, the concept of knowledge primarily functions as a guideline for whether people's actions should be blamed or praised.

[150] Ibid., p. 9.

Of course, Hawthorne and Stanley do not actually perform a four step conceptual synthesis. Rather, they begin straightaway with intuitions about the concept of knowledge by listing several everyday situations in which it intuitively seems appropriate for one subject to appraise another's behavior by using the concept of knowledge.

> … our ordinary folk appraisals of the behavior of others suggest that the concept of knowledge is intimately intertwined with the rationality of action.[151]

For example, in their version of "Selling a Lottery Ticket," they present a subject who does decide to sell his ticket, using "I will lose the lottery" as his reason for action.[152] They go on to encourage the reader to notice the absurdity of such reasoning, that is, to appraise his behavior as irrational. In another example John decides not to buy health insurance.[153] His mother then appraises his behavior as irrational, due to the fact that John doesn't know that he won't fall ill. Other examples include conditional orders, "Don't take out the cake until it's done" implies "Don't take out the cake until you know it's done," and judgments of negligence, "Only a negligent doctor would use a needle that he doesn't know is safe."[154] Noting the way one intuitively appraises such behavior, they offer the Reason-Knowledge Principle as the foundation of these normative judgments. Namely, one tends to criticize subjects who use a premise that they do not know as a reason for action. "Folk appraisals" of the kind of behavior these subjects exemplify motivate the Reason-Knowledge Principle.

KAP on the other hand is motivated by scenarios in which inquirers must decide whether or not an informant's assertion is trustworthy and scenarios in which inquirers must decide whether or not they may acceptably stop inquiry. The rationality of an action is thus taken as an independently determined point.[155] Of course, a knowledge claim can be used to evaluate behavior; however, the accuracy of a knowledge claim ultimately rests on the rationality of the action, not the other way around. Therefore, on my hypothesis, the concept of knowledge does not enable

[151] Ibid., p. 1.
[152] Ibid., pp. 1 – 2.
[153] Hawthorne & Stanley 2008, p. 1.
[154] Ibid., pp. 2 – 3.
[155] Accordingly, my account would not be vulnerable to Igor Douven's objection that rational practical reasoning is best represented with a Bayesian model and does not demand the concept of knowledge. (Douven, 2008)

us to evaluate action; rather, its sensitivity to action evaluations enables us to evaluate beliefs. So, while KAP could be seen as motivated by evaluation, the primary object of evaluation would be a belief or an assertion, not an action.

iii. First person scenarios

The Reason-Knowledge Principle describes the relationship between a practical reasoner and a proposition. It says that a subject should not use a proposition as a reason for action unless she herself knows it.

KAP is more flexible. It too can describe cases in which a subject may use a proposition as a reason for action because she herself knows it. However, it also describes how the concept of knowledge functions in cases in which a subject can use a proposition as a reason for action because some other, external, informant knows it. As we saw in the examples above, this flexibility led KAP to entail a contextualist account of knowledge. The police may truly say that "Louise does not know that John was at work," thanks to their practical environment. Thelma may truly say that "Louise does know that John was at work," thanks to her practical environment.

The Reason-Knowledge Principle does not entail contextualism. Rather, it entails subject-sensitive invariantism. Given the driving hypothesis, there is no reason to explicate sensitivity to the practical environments of hypothetical inquirers. Only the subject's (S_1's) own practical environment is relevant to whether or not he knows. Hawthorne introduces this approach as follows:

> Restricting ourselves to extensional matters, the verb 'know' picks out the same ordered triples of subject, time, and proposition in the mouths of any ascriber. However, whether a particular subject-time-proposition triple is included in the extension of 'know' depends not merely upon the kinds of factors traditionally adverted to in accounts of knowledge – whether the subject believes the proposition, whether that proposition is true, whether the subject has good evidence, whether the subject is using a reliable method, and so on – but also upon the kinds of factors that in the contextualist's hands make for ascriber-dependence. These factors will thus include (some or all of) the attention, interests, and stakes of that subject at that time.[156]

This approach to the relationship between knowledge attributions and practical considerations is called subject-sensitive invariantism. Knowledge is sensitive to the practical concerns of the subject, but not to those of the attributor. For example, say

[156] Hawthorne 2004, p. 158.

that S_1's stakes are high in the morning and low in the evening, but otherwise his relationship with *p* remains identical – he believes it and has some set amount of justification. It would then be possible for S_1 not to know *p* in the morning, but to know p in the evening. If an attributor has low stakes the entire time, then her assertion that 'S_1 knows that p' would be false in the morning and true in the evening. According to subject-sensitive invariantism, the context of an attributor is irrelevant to the truth of her attribution.

Subject-sensitive invariantism would describe Louise's epistemic status in the John example as follows: As long as Louise is with Thelma, governed by the standards of the petty bet practical environment, then Louise does know that John was at work. As long as Louise is with the police, she does not know. If the police assume that Louise does not know based on the strict demands of their *own* practical environment, then that is simply a mistake of over-generalization; they accidentally apply their own standards to Louise.[157] The police cannot make any accurate claims about whether or not Louise knows without being privy to Louise's practical environment, because the concept of knowledge is sensitive only to the practical environment of the knower herself. Thus, if Louise is with Thelma in the petty bet environment, then the police *cannot* truly say that "Louise doesn't know whether John was at work."

Since the hypothesis driving the Reason-Knowledge Principle only considers the relationship between one subject's reasoning and *her own* belief, we may only explicate a first-person sensitivity from it, according to which the reliability condition varies exclusively with *the knower's own* practical environment. Since KAP is based on the hypothesis that the concept of knowledge may also function in the relationship between one subject's reasoning and a *different* subject's belief, it goes on to explicate a reliability condition that may also vary with the *attributor's* practical environment. Thus the Reason-Knowledge Principle entails subject-sensitive invariantism, while KAP entails contextualism.

iv. Knowledge per se

Hawthorne and Stanley claim that the truth of knowledge attributions is only sensitive to the subject's practical environment. This makes their account a form of

[157] Ibid., pp. 164 – 165. For a convincing objection to this error theory, see Cohen 2004, p. 489.

invariantism, for the meaning of a knowledge attribution cannot change according to the attributor's concerns. One subject can only have one practical environment at a time.[158] There is no such limit to the amount of attributors that may be talking about a certain subject. Furthermore, each attributor might have her own unique practical environment. As we saw above, that fact leads KAP to endorse the contextualist claim that "S knows that *p*" and "S doesn't know that *p*" might simultaneously be true in the mouths of two different attributors. Therefore, if KAP is correct, we will sometimes be unable to offer a final diagnosis of a given subject's epistemic position, for it is always possible that an attributor's practical environment will change.[159] Subject-sensitive invariantism does not lead to this situation. That is, it does assume that, for a given subject at a given point in time, there is a conclusive answer to the question "Does S know that *p*?" At that point in time, S either knows or he does not know. Therefore, according to subject-sensitive invariantism, we are not limited to judging whether S or someone else could truthfully say "S knows," but may also judge whether or not S knows *per se*.[160]

v. How to compare the two kinds of sensitivity

Most work on comparing contextualism to subject-sensitive invariantism has relied on examining intuitions about the appropriateness of different sentences. As explained in section 4.b.iii, this method of examination is not as well suited to a conceptual synthesis. Accordingly, we must judge the respective merits of the two accounts with different tools.

A practitioner of conceptual synthesis should ask which theory can better account for the function of knowledge attributions. KAP is motivated by the way knowledge attributions function in evaluating beliefs and attributions. Subject-sensitive invariantism is motivated by the way knowledge attributions function in evaluating behavior. In order to compare the two theories, we should ask whether one of the theories can better account for both functions.

---

[158] Although, contrast the example of Jane, the walking talker in DeRose 2009, pp. 270 – 271. I return to this example in 6.b.

[159] Sometimes a contextualist *can* offer a final diagnosis. For example, a contextualist could agree that "S knows that *p*" is conclusively false if *p* is false.

[160] As Michael Williams would say, subject-sensitive invariantism, but not contextualism, is compatible with epistemological realism. See Williams M. 1996, Ch. 3 and Williams M 2001, pp. 193 – 194.

I will argue in section 6.b that the answer to that question suggests the superiority of KAP to subject-sensitive invariantism. That is, KAP can account for the role knowledge plays in behavior appraisal, while a subject-sensitive invariant synthesis cannot account for the role knowledge plays in third-person belief appraisal.

d. Two objections to the private nature of KAP

By hypothesizing that the concept of knowledge plays a role in practical reasoning, I have explicated a contextually variant reliability condition. However, in so doing, I have ignored the social circumstances in which any system of epistemic evaluation must be embedded. The time has come to confront the tension between the social need for a stable category and the personal need for sensitivity to practical environments. If my claim is correct and the concept of knowledge can fill both needs, then we need some explanation of how this is possible. I have explicated two seemingly contradictory conditions and must now explore whether they may be synthesized into one concept. The following two objections demonstrate the difficulty of such a synthesis.

i. The authority objection

(1) Craig: recommendations/flags

Craig emphasizes that people do not only notice informants who are good for them; they also help each other find good informants. He points out that the concept of knowledge has adapted to this social practice; it is not indexed to the private needs of one person, but rather suited for use in *recommendations*. Thus, one role for knowledge attributions lies in cases of recommendation. When an inquirer attributes knowledge to an informant, she not only trusts him as to whether *p*, she also declares that he is someone who *should* be trusted as to whether *p*, thereby flagging him and recommending him to anyone else who might need to find out whether *p*.

> To put it briefly and roughly, the concept of knowledge is used to flag approved sources of information.[161]

[161] Craig 1990, p. 11.

KAP does not suffice to explain how a knowledge attribution can function as a flag of epistemic authority. KAP shows that the truth of a knowledge attribution may be solely dependent on whether S_1 would have authority *over* S_2. If Estelle could act on Nelville's authority that *p*, then Estelle can truthfully say that "Nelville knows that *p*." However, a flag is by definition meant to be used by others. If the spoken knowledge attribution is to be able to function as a flag, the truth of the knowledge attribution must also depend on whether S_1 has authority over the person *hearing* the attribution, not only the person *making* it. That is, we must consider some S_3, who is hearing or who may overhear the attribution. If S_1 is to have appropriate epistemic authority over S_3, S_3 must be able to use *p* as premise in rational practical reasoning. If the concept of knowledge is to function properly in recommendations, we need to be able to say that Estelle would be *wrong* to tell Ida that "Nelville knows that the cookies have no peanuts." Yet KAP claims that a knowledge attribution can depend exclusively on Estelle's personal needs; we have as yet explicated no condition about what happens when an S_2 meets an S_3.

If a knowledge attribution is to function in recommendations, then it must be a flag of *public* epistemic authority. Yet, contextualists typically conceive of practical environments as a *private* affair.[162] The fact that S_1's assertion is good enough to use in S_2's practical environment offers no guarantee that it will be good enough to use in anyone else's practical environment, such as that of S_3. Therefore, it seems that we cannot describe the concept of knowledge as relative to an attributor's practical environment if we also claim that it plays a functional role in recommendations.

(2) Epistemic authority confers an honorable social status with normative implications

In Chapter Three I suggested that we substitute talk of 'objectively good informants' with talk of 'epistemic authority'. I supported this substitution by pointing to intuitive connections between the concept of authority and the social conditions that we will have to explicate, if the synthesized concept is to match up well with knowledge. These conditions go beyond Craig's need for recommendations; I argued that people also need the concept of epistemic authority to imply a social status, honor, and a normative aspect. All three of these points seem to directly conflict with

[162] See section 5.c.

the contextual sensitivity explicated above. KAP entails that a subject's knowledge can come and go relative to which hypothetical relationship one has in mind.[163] If Olivia's practical environment implies extremely low standards, she may be able to truthfully apply knowledge to Arlo, despite the fact that hardly anyone else would be able to act on Arlo's authority. Given this fickle variability, a fixed social status for knowers seems out of the question. There would be equally little reason to honor someone for such a transient feature. If "knows" can be as worthless a predicate as it is in Olivia's mouth, people would not intuitively connect it with value. Finally, the contextual sensitivity threatens the normative implications as well. Olivia can truthfully say that Arlo knows, despite the fact that others would be irrational to believe him. The normative condition that people *ought to believe* a knower conflicts with a contextually variant reliability condition.

ii. The intelligibility objection

Craig argues that an undetermined group of potential inquirers must be able to trust on the authority of a knowledge attribution. However, there is another requirement which must be filled before the question of authority can even be broached, an intelligibility requirement. If inquirers cannot even understand the spoken attribution, then they will not be able to trust on its basis either.

(1) Disquotation

Contextualist accounts make intelligibility seem mysterious. If the meaning of a word is always shifting according to the context of the speaker, how can an inquirer understand what is being communicated by a knowledge attribution? Of course, many intelligible words mean different things in different contexts, such as 'I', 'tall', or 'nearby'. However, some object that our linguistic intuitions do not treat 'knowledge' as they do those other words. In particular, those other words are not disquotable, whereas knowledge attributions are disquotable.[164] If a word is disquotable, then I can report someone's use of it without using quotation marks, i.e. without explicitly drawing attention to the (possibly different) context she was in

[163] That is, the truth value of "S knows that *p*" can change according to the hypothetical relationship the speaker has in mind. If you see an important difference between the two phrases, please reference DeRose's discussion of the fallacy of semantic descent. (DeRose, 2009, p. 218)
[164] Hawthorne 2004, pp. 98 – 111. Davis 2007, pp. 399 – 401.

when she said it. Words that change their meaning with context are not disquotable, whereas words with invariant meanings are. 'Nearby' is not disquotable. Janice might tell me on the phone, "Jack is nearby." It would be absurd for me, standing elsewhere, to disquote her report by claiming "Janice believes that Jack is nearby." Such a statement would be false. Rather, if I repeat her statement, I would have to use the word 'nearby' in quotes in order to indicate that it was true for her, the original speaker, not for me: "Janice believes that Jack is 'nearby' – that is, near the tree, where she is."

If 'knowledge' varied contextually, then it too would not be disquotable. Imagine that Estelle sincerely asserts that "Nelville knows that *p*." Imagine further, that Ida overhears that sincere assertion. If the truth values of knowledge attributions vary contextually, then Ida's assertion that "Estelle believes that Nelville knows that *p*" is not true, because Ida's practical environment is different than Estelle's practical environment. Thus, "knows" in Estelle's mouth means something different than "knows" in Ida's mouth. Ida could only ensure accuracy by saying something like, "Estelle believes that Nelville 'knows' that *p*, that is, knows according to the standards of knowledge implied by her practical environment." Only with a modifying clause may Ida disquote the word "know." Therefore, if contextualism is correct, people would have to report knowledge attributions in the same way they report nearness attributions, explicitly stating the situation in which they were made.

However, according to Hawthorne, people do regularly disquote knowledge attributions without adding such modifying clauses. That is, in most cases, S_3 will simply claim that "S_2 believes that S_1 knows that *p*" without making any reference to the practical environment S_2 was in when she made her attribution.

Of course, this objection depends on the way we actually use knowledge attributions, which brings us prematurely to the fourth comparative step. Nonetheless, the point does affect the ability of a concept to play the hypothesized function, making it appropriate to consider now during the third explicative step. If a concept is to function in a social system of epistemic evaluation, it must be intelligible. If we, in fact, disquote knowledge attributions without adding modifying clauses, then the concept of knowledge, given my explication, would be unintelligible. So, we have the following conflict: (a), according to the explication, the meaning of 'knowledge' varies according to the original attributor's context,

while (b), according to empirical fact, speakers usually report knowledge attributions without referring to the attributor's context (i.e. without quotes). It seems that in order to save the explication, I would have to hold that our usual, disquoting reports of knowledge attributions are usually false. Yet, if we systematically misuse or misunderstand the word 'knowledge' in these situations, then the concept cannot be properly functioning as a term in a social system of epistemic evaluation. Therefore, I must somehow explain how knowledge can vary contextually without opposing our disquotation related intuitions.

(2) Bach: Eavesdropping

Kent Bach provides a vivid example of the intersubjective intelligibility of knowledge attributions that does not explicitly rely on disquotability, though it exploits the same tension. He imagines a case in which someone eavesdrops on a knowledge attribution, and thus presumably would be utterly clueless as to what kind of practical environment the speaker meant to imply.

> What does contextualism predict if you encounter a [knowledge] attribution out of context? It seems to predict that you won't be in a position to grasp which proposition the sentence expresses. Suppose you eavesdrop on the middle of a conversation and hear one person say to the other, "Nixon knew that Liddy was planning the Watergate break-in." Since it is not evident to you which [knowledge] relation 'knew' expresses, you can have only a vague idea of what is being said. Lacking any specific information about the context in which the [knowledge] attribution was made, you should feel a bit uncertain as to what was said. But you won't.[165]

With the eavesdropping case, Bach demonstrates that knowledge attributions do not require explicit clarification. An eavesdropper can understand what a speaker is implying with his knowledge attribution, even if the speaker isn't talking to her at all.

(3) Memory

My presentation of the intelligibility and authority objections has thus far focused on the difficulties of intersubjective communication. However, difficulties can also arise for one person at two different points in time. In other words, if a person makes a knowledge attribution, spoken or unspoken, and then later remembers it, how is it

[165] Bach 2005, p. 61.

possible for that attribution to be authoritative and intelligible?[166] A functional analysis must allow for a knowledge attribution to temporally endure.

e. Three desiderata

I have argued that the concept of knowledge serves a function in practical reasoning: by finding people who know, an inquirer finds people with the epistemic authority necessary to provide her with premises she can rationally use in her practical reasoning. I have argued further that the concept of knowledge serves a function by flagging that epistemic authority for the public.

Therefore, if this functional synthesis is to succeed, we must find some way to synthesize the following three conditions into one coherent concept:

(1) Sensitivity to practical environment

In order for "knowledge" to function as a flag of epistemic authority, its meaning must vary with the features of the practical environment in which it is attributed.

(2) Intersubjective authority

In order for "knowledge" to function as a flag of epistemic authority, it must imply epistemic authority to listeners.

(3) Intersubjective intelligibility

In order for "knowledge" to function as a flag of epistemic authority, it must be intelligible to listeners.

In the next chapter I evaluate several accounts, to see how they fare at satisfying these three desiderata. If one account can simultaneously satisfy all three, then it is indeed possible for these seemingly contradictory conditions to be synthesized together into one coherent concept.

[166] Hawthorne 2004, p. 111. Williamson 2005, p. 101.

5. In Search of a Synthesis

The objections raised at the end of Chapter Four threaten to disprove the hypothesis by making a coherent synthesis impossible. If no single coherent concept could possibly meet the need for epistemic authority, then the concept of knowledge must have some other function. Thus, while Craig moves seamlessly from explication to synthesis, this step poses a significant problem for my own account.

In the following chapter, I shall examine several conceptions of knowledge, to judge whether they might provide a model that can indeed satisfy all three desiderata. First I outline two invariantist conceptions of knowledge: Craig's high standards account and Bach's moderate account. Neither provides a synthesis that would meet the need for sensitivity. Next I outline two portrayals of contextualism: salience-sensitive contextualism and pragmatic contextualism. Neither explains how knowledge attributions can be intersubjectively authoritative and intelligible. Finally, I introduce the relatively novel solution of convention contextualism: construe contextual variation as sensitive only to *publically recognized* contexts. I argue that this option can satisfy all three desiderata, thereby demonstrating that my explicated conditions can indeed come together in a single coherent synthesis.

a. High standards invariantism

In his final objectivized synthesis, Craig leaves no room for contextual sensitivity.[167] He argues that, to accommodate a community of users, the concept of knowledge must be reliable enough to be authoritative for any individual, no matter how stringent the demands of her practical environment. I have already presented his motivation for this stance,[168] so here a brief summary will suffice. A knowledge attribution made in one practical environment must be able to flag a good informant for a second inquirer in a different practical environment. That is, one must be able to report that "S_1 knows that *p*" across contexts, without risking an unintentional falsehood. We are not always sure what sort of practical environment our interlocutors are in, or to whom they will report our attribution. Thus the failure to function in trans-contextual reports would lead to the failure to function in any kind

[167] In the appendix to his book, Craig does suggest that some contextual sensitivity might be possible. To that extent he supports Unger's semantic relativity, which I discuss below in section d.
[168] See 2.a and 4.d.

of intersubjective reports. Basically, if its reliability were indexed to the speaker's personal concerns, the concept of knowledge would not be able to serve any sort of public function in spoken attributions.

Craig argues that the only way a knowledge attribution can serve its public function is by adapting to the demands of the most stringent practical environment. This adaptation puts Craig's objectivized concept in accordance with high standards invariantism.[169] That is, "S_1 knows that *p*" must consistently imply such a high likelihood of *p* being true that *anyone* in *any* situation could use *p* as a premise in rational practical reasoning.

> [Objectivization] is going to edge us towards the idea of someone who is a good informant as to whether p whatever the particular circumstances of the inquirer, whatever rewards and penalties hang over him and whatever his attitude to them. That means someone with a very high degree of reliability, someone who is very likely to be right – for he must be acceptable even to a very demanding inquirer. So of the worlds that we cannot quite definitely exclude, we shall want to include in our assessment of him even those that we regard as very improbable. Moreover, we shall be motivated to take a pretty careful look at those which we 'can quite definitely exclude' – is that really as many as we think? These thoughts take us further down the road of objectivisation. Knowledge, so the hypothesis goes, lies at the end of it.[170]

Unsurprisingly, this high standards invariantism does well on the authority and intelligibility desiderata. High standards *invariantism* ensures that a knowledge attribution can only mean one thing, and would therefore always be intelligible, regardless of the listener's amount of background information. *High standards* invariantism holds that a knowledge attribution expresses absolute epistemic authority in any practical environment. If knowledge were to imply such infallible authority, there would be no context in which a subject would not be able to use the attribution as the mark of an epistemically authoritative informant.[171]

Nonetheless, Craig's high standards invariant account conflicts with the hypothesis that the concept of knowledge plays a function related to practical reasoning. On his account, the concept of knowledge could not be used to recommend good informants without risking falsehood. While using "knowledge" to

[169] Kelp points out that Craig's characterization of objective knowledge is best described as high standards invariantism. (Kelp 2001, p. 7.)
[170] Craig 1990, p. 91.
[171] See Fantl & McGrath 2002, pp. 84 – 87 for a discussion and refutation of the analogous position that "S is justified in believing that p only if anyone with S's evidence for p, no matter what the stakes, would be rational to prefer [and hence act] as if p."

recommend good informants could never be misunderstood, using the *denial* of "knowledge" to advise against bad informants could very often be misunderstood. There would be a plethora of situations in which S_2 could truly tell S_3, "S_1 does not know that *p*," despite the fact that S_3 would in fact be able to use *p* as a premise in rational practical reasoning. If high strung Hannah tells sensible Sue not to trust Albie the Average Informant, then on Craig's account, Hannah is telling the truth, yet Sue would be irrational to believe her.[172] Only when the standards of a knowledge attribution are, for both recommender and inquirer, in all respects as strict as the most demanding practical environment would ever require, would the concept of knowledge reflect the relevant standards of epistemic authority. In all ordinary situations, people would be acting entirely rationally to use *p* as a reason for action, despite believing that *p* on the authority of a subject who did not know that *p*.[173] Hence, the concept of knowledge would not serve a function in practical reasoning - not the function implied by Craig's prototypical scenario, and not the function implied by the recommendation scenario either. The concept of knowledge would be identifiable with neither the concept of the good informant nor the concept of the objectively good informant.

In order to serve the hypothesized function, the concept of knowledge must be sensitive enough to practical reasoning so as not to be irrelevant to ordinary concerns, yet insensitive enough so as not to be insufficient for extraordinary concerns. By adapting his synthesis exclusively to the latter condition, Craig manages to satisfy the authority condition at the cost of the sensitivity condition. Yet, as I have presented the explication, our synthesized concept demands both.

Verdict

High standards invariantism can account for the authority and intelligibility of the concept of knowledge, but only at the cost of sensitivity.

[172] Feldman makes the same argument, concluding that there is no good reason why objectivization should result in such high standards for good informants. (Feldman, 1997, p. 211)
[173] Fantl & McGrath 2002, p. 85.

b. Moderate invariantism

i. Misleading intuitions: Bach's error theory

As discussed above, functional contextualism can be satisfied by either pragmatic or semantic sensitivity to practical environment. However, Kent Bach has argued that the concept of knowledge exhibits neither sort of sensitivity; rather, it is semantically and pragmatically invariant.[174] According to his analysis, a knowledge attribution always implies the same amount of authority, but that amount is set reachably low. Thus, he defends a form of moderate invariantism.

The examples I presented above in section 4.b showed situations in which the concept of knowledge does exhibit sensitivity to practical environments. It is generally accepted that these examples are intuitively acceptable; we tend to approve of the way in which knowledge is used in the given cases.[175] Since Bach can resort to neither semantic nor pragmatic concerns to explain these intuitions, he must offer some kind of error theory.[176] That is, he must argue that we make a mistake when we intuitively accept the contextual sensitivity of knowledge attributions demonstrated in the given examples.

He begins with a point of general agreement: "knowing that p requires that one's experience/evidence/justification rule out counterpossibilities (alternatives to p, threats to the basis for one's belief that p)."[177] In other words, knowledge requires the ruling out of relevant alternatives.[178] Now we are left with the question of *which* counterpossibilities must be ruled out in order for the belief to count as knowledge. If we insist that any conceivable counterpossibility must be ruled out, we are left with radical skepticism. Bach rejects skepticism and thereby argues that one need rule out only *relevant* possibilities. While he does not present any specific criteria for judging whether an alternative is relevant, he does maintain that relevance can be determined by evidential concerns alone. He thereby rejects both contextualism and subject-sensitive invariantism, for he denies that the proper definition of relevance involves anyone's practical environment.

[174] Bach 2004. Bach 2005.
[175] I am assuming that my reformulations of the examples from their source texts have not affected their intuitive acceptability.
[176] Bach 2005, p. 78.
[177] Ibid., p. 81.
[178] See above, section 4.b.i.

The next step is to explain why we intuitively take a practical environment as affecting the relevance of a counterpossibility; if that is just a mistake, why is it so widespread? Bach begins his answer to that question by pointing out that "somehow, at least to some extent, our belief-forming processes are tuned into plausible sources of error."[179] As long as a subject is healthy, unafflicted by paranoia, hypochondria or some other disorder, then irrelevant counterpossibilities do not regularly occur to her. Accordingly, the fact that a counterpossibility to some proposition occurs to a subject at all is usually itself a piece of evidence that the counterpossibility is epistemically relevant. The reverse also holds; the fact that a counterpossibility to some proposition does *not* occur to a subject usually constitutes evidence that the counterpossibility is *not* epistemically relevant. Bach offers the airport example, in which Mary is wondering if her flight will stop in Chicago. If it occurs to Mary that her flight itinerary might have a misprint, then the very fact that this counterpossibility occurs to her is usually a valuable piece of evidence, suggesting that she must in fact attempt to find some more reliable source if she wants to really know whether her flight will stop in Chicago. The general reliability of our belief-forming processes means that we can take our own regard or disregard of a counterpossibility as valuable evidence of that counterpossibility's relevance.

Bach stresses, though, that such evidence is highly defeasible. A great deal of factors can cause a subject to overlook a relevant counterpossibility or attend to an irrelevant counterpossibility. Mental illness is one such factor: if Mary is a paranoid schizophrenic, then the fact that the possibility of a misprint on her flight itinerary occurs to her is *not* evidence for the relevance of that possibility. Bach argues that high stakes can be another of those merely misleading factors. If Mary's life depends on whether her flight stops in Chicago, the possibility of a misprint is more likely to occur to her, despite the fact that it is *actually* irrelevant. High stakes can cause epistemically irrelevant counterpossibilities to occur to a subject.

With this set-up, Bach has the tools to explain why our intuitions are likely to mislead us when we evaluate examples that involve high stakes. Since the occurrence of a live counterpossibility to *p* is *usually* evidence that the subject doesn't know that *p*, Mary mistakenly accepts it as evidence in the high stakes airport case, and mistakenly concludes that she does not know that the flight will stop in Chicago. We

[179] Ibid., p. 84.

who read the example then make the same mistake and intuitively agree with Mary's conclusions that she does not know. We accidentally rely on evidence which is usually reliable, but in this particular case is misleading.

Thus, a knowledge attribution does not always retain authority on the moderate invariantist account. In situations such as Mary's, practical standards can be so high that knowledge does not imply as much justification as the practical reasoner requires. A knowledge attribution only carries enough epistemic authority for *average* situations. In particularly demanding practical environments, Smith might know that *p,* and yet his relationship to *p* might still not be good enough for Mary, with her life and death stakes, to believe that *p* on his authority. As Bach puts it, "Sometimes it is reasonable to go beyond the call of epistemic duty."[180]

ii. Belief removal

We cannot expect any invariantist account to fare well on the sensitivity desideratum. Nonetheless, Bach might argue that his account can indeed accommodate the intuition that a subject should only use *p* as a reason for action, as long as it is on the authority of someone's knowledge that *p*. His tactic to accommodate this intuition involves the phenomenon of *belief removal*. He points out that high stakes generally cause the subject to stop believing *p* herself.[181] That is, they raise her doxastic threshold, making her demand more evidence before she will commit to a belief. Since Mary no longer *believes* that her flight stops in Chicago, it is true that she no longer *knows* it either, since belief is necessary for knowledge. Therefore, the intuition that she should not use "the flight stops in Chicago" as a reason for action is correct and further, does not conflict with KAP's claim that knowing that *p* ensures that one may rationally use *p* as a reason for action. Mary cannot use *p* as a reason for action, but neither does she know that *p*.

Jessica Brown (2005) argues that Bach's view is nonetheless an unsuccessful accommodation of our practical reasoning related intuitions in high stakes cases. It relies on a "brute contingent psychological fact,"[182] namely that the subject in the high stakes scenario does stop believing that *p*. If the subject dogmatically continues to believe that *p*, then, according to Bach, she knows that *p*. If Mary, despite her

[180] Bach 2004, p. 14.
[181] Bach 2005, p. 77.
[182] Brown 2005, p. 84.

astronomical stakes, continues to believe that the flight will stop in Chicago and refuses to consider the alternative of a misprint, then Bach must admit that she knows that *p*. Yet, she still cannot use *p* as a premise in rational practical reasoning. Thus, Bach's belief-removal maneuver cannot change the inability of his account to accommodate intuitions about the connection between knowledge and practical reasoning.[183]

With or without belief-removal, Bach cannot accommodate intuitions about the appropriateness of *third-person* knowledge attributions. According to Bach, Mary's inability in a high stakes situation to believe that *p* does not affect *Smith's* knowledge that *p*. Since Mary does not believe that the airplane will stop in Chicago, her claim that "I do not know that the airplane will stop in Chicago" is true. Nonetheless, her claim that "Smith does not know that the airplane will stop in Chicago" is false. Since KAP is based on the hypothesis that the concept of knowledge functions as a flag of third-person epistemic authority, moderate invariantism fails by KAP's standards of sensitivity, despite the belief removal maneuver.

iii. Objection to the arbitrariness of semantic error theories

We have seen that moderate invariantism fully accounts for neither sensitivity nor authority. That is not the only objection worth offering to moderate invariantism. Bach stresses that his account is a semantic theory, not a substantial epistemological theory of knowledge. He therefore doesn't feel obliged to offer a precise answer as to just what may count as a relevant alternative. However, we have seen that his account does attribute certain systematic errors to speakers. If his account is merely meant to *reflect* semantic use without postulating an epistemological theory, then what criterion can he have for *criticizing* the judgments of subjects in high stakes practical environments, calling them errors?[184]

One may argue that semantic accounts should strive for attributing the fewest errors to speakers and general intuitions; similarly, the best account would accommodate the most intuitions. However, it is not a matter of common consent

[183] Brown also points to Bach's concomitant inability to account for common intuitions as to the appropriateness of a self-attribution of knowledge that *p* and the assertion that *p*.
[184] See Owens 2000, p. 26 for the argument that no purely evidentialist account of justification has the resources to specify exactly what level of evidence is demanded by knowledge; only by referring to the believer's needs and interests could we possibly provide a specific level.

that speakers do make only a minimal amount of errors when attributing knowledge. Wayne Davis, for example, argues that people regularly misuse the word 'knowledge'.[185] By adhering to the standard of attributing the fewest errors, semantic theorists betray substantial epistemological convictions. Those convictions must be defended and it is difficult to discern what tools Bach would have to do so.

KAP does not run into this problem. On the strength of my functional hypothesis, I offer a standard by which to determine the relevance of a counterpossibility: namely, if it affects the acceptability of that premise in rational practical reasoning, then it is relevant and must be ruled out before knowledge can correctly be claimed. That standard is not arbitrary, but rather motivated by the hypothesis that societies need the concept of knowledge to function in practical reasoning.

Verdict

As with all invariant accounts, intelligibility is no mystery. Yet moderate invariantism can fully account for neither the sensitivity nor the authority of knowledge attributions. Furthermore, the error theory which moderate invariantism must provide is inadequately motivated.

c. Salience-sensitive contextualism

i. The Rule of Attention

Some versions of contextualism allow for a very great deal of fluctuation in what is implied by a knowledge attribution. 'Salience-sensitive contextualism' refers to those particularly volatile accounts which maintain that salience to an individual subject is enough to make an alternative relevant.

Remember, relevant alternatives are possibilities that a subject must be able to rule out if he is to count as knowing that *p*. For example, it may be that I cannot know that the light is on if I cannot rule out the alternative that you have turned it off. David Lewis offers several rules which may be used to determine whether or not an alternative is relevant, i.e. whether it must be ruled out. One of his rules is the Rule of Attention, according to which whatever alternatives a certain speaker is paying attention to are automatically relevant.

[185] Davis 2007, p. 428.

> What is and what is not being ignored is a feature of the particular conversational context. No matter how far-fetched a certain possibility may be, no matter how properly we might have ignored it in some other context, if in *this* context we are not in fact ignoring it but attending to it, then for us now it is a relevant alternative. It is in the contextually determined domain.[186]

That is, a speaker's context is in part determined by the *salience* of possibilities to that speaker. The meaning of her knowledge attribution changes relative to whatever she happens to be thinking about. One moment 'knows' might mean "able to rule out alternatives 1, 2, and 3" while moments later, it might mean "able to rule out alternatives 1, 2, 3, 4, ... 25" thanks only to a host of random paranoid fantasies suddenly occurring to the speaker. This leaves the concept of knowledge with a very high level of sensitivity, but makes intersubjective authority and intelligibility extremely puzzling.[187]

We may better understand the appeal of the Rule of Attention when we consider its historical motivation. Lewis's account is not based on any hypothesis about the function of knowledge attributions, but rather an attempt to offer a satisfying diagnosis of the power of skeptical arguments.[188] When confronted with the possibility that everything we believe is based on mere illusion, and our inability to eliminate that possibility, people usually react by accepting the skeptical conclusion that we do not in fact know anything at all.[189] Salience-sensitive contextualists can explain that phenomenon by arguing that the salience of a possibility makes the possibility relevant to the thinker, so that, thanks to the change in context, the skeptical conclusion is actually true relative to her standards. Only once the skeptical possibilities have been forgotten, do the everyday claims to knowledge become true again.

My synthesis is not motivated by the desire to interpret the allure of skepticism. In fact, I have narrowed my focus to only those situations in which knowledge plays a role in practical reasoning; skeptical arguments have very little sway there. With functional rather than skeptical concerns at the forefront, the Rule of Attention does

---

[186] Lewis 1996, p. 559.

[187] See Hawthorne 2004, pp. 158 - 173 for a different discussion and refutation of salience contextualism, based on the observation that the salience of counterpossibilities most plausibly affects knowledge by means of belief removal, which can be accounted for by a traditional invariantist account.

[188] Lewis 1996, p. 549.

[189] As Michael Williams points out, *accepting* the skeptical argument is the obvious and transparent move. It is the rejections of skeptical arguments that involve arcane philosophical theories. (Williams M. 1996, p. xix)

more harm than good, as it is difficult to imagine how such a wildly fluctuating concept could command the authority and intelligibility required to properly function.

ii. DeRose's rejection of the intelligibility objection

Keith DeRose follows Lewis in many respects, including an acceptance of the Rule of Attention and the rapid fluctuation such a rule entails. DeRose maintains that a speaker's own standards at a given point in time fully determine the meaning of the word 'knowledge' in her mouth, no matter how unreasonable those standards may be.[190] However, DeRose goes on to offer a contextualist response to the intelligibility problem: namely, he denies that it is a problem. Remember, the objection was based in part on a supposedly empirical fact: people regularly *do* in fact disquote knowledge attributions without modifying clauses. DeRose denies that linguistic data show any such thing; when 'knowledge' is disquoted, it is not across contexts. He thus accepts that "knowledge" is not necessarily intelligible across contexts.[191] When the hearer and the speaker's practical environments differ enough, then the hearer cannot and would not disquote the speaker's attribution.

However, DeRose goes on to admit that "knowledge" is sometimes disquoted.[192] If one overheard a conversation in which Frank said, "Mary knows that the library closes at 5" then one would not hesitate to disquote and say, "Frank believes that Mary knows that the library closes at 5." He diffuses the objection by comparing the word "knowledge" to the word "tall," pointing out that "tall" is just as often disquotable. [193] If one eavesdropped on a conversation in which Frank said, "Mary is tall," one would not hesitate to disquote and say, "Frank believes that Mary is tall." Yet it is uncontroversial that "tall" is context sensitive. Rather than go on to offer a positive explanation as to how and why this common practice avoids massive confusion, DeRose contents himself with a challenge: anyone who finds the intelligibility problem a threat to contextualism about "knowledge" must also demonstrate why it is not a threat to contextualism about "tall."

---

[190] DeRose 2009, p. 142.
[191] Ibid., p. 165. See Greco 2008, p. 434 for the same argument.
[192] DeRose 2009, p. 167.
[193] Ibid., pp. 169 – 170.

However, DeRose neglects to address an important difference between the words "tall" and "knows." A knowledge attribution, on the functional account, must operate as a flag of authority. Whether or not someone performs a certain action can depend on whether or not a knowledge attribution is true or false. Of course, while tallness is not involved in any general norm of action, the truth of a tallness attribution can still occasionally affect someone's actions. However, I presume that those are precisely the instances in which "tall" is no longer dequotable across contexts. That is, if Lucius will only pick Reeves for the basketball team on the authority of Theo's assertion that "Reeves is tall" then Lucius will need to acquire enough background information to ensure that he and Theo share the same practical environment. The problem is that, according to KAP, a knowledge attribution must almost *always* be authoritative enough to imply that one could act on its basis. In that case, an inquirer would almost always need to ensure that she shared a practical environment with the speaker of the knowledge attribution before she could acceptably disquote the attribution.

"Know(s)" must almost always imply epistemic authority, whereas "tall" can often lack the level of authority needed to motivate action. However, given DeRose's claim that knowledge is sensitive to the salience of alternatives to an individual, his account entails that the meaning of a knowledge attribution is likely to change radically across contexts. We would then expect that people would be more cautious in disquoting 'know(s)' than they would be when disquoting 'tall'. Thanks to the epistemic authority condition, there is a reason to think that the intelligibility problem is threatening to contextualism about "knowledge" despite the fact that it poses no threat to contextualism about "tall."

Verdict

Salience contextualism allows for a great deal of sensitivity, but pays the price by failing to satisfy the intelligibility and authority desiderata. DeRose rejects the intelligibility objection, claiming that if it were a problem for "knowledge," then it would be a problem for "tall" too. However the additional demands of the authority desideratum suggest that the intelligibility objection does have more force for "knowledge" than it does for "tall".

d. Pragmatic contextualism

i. Semantic relativity

Although passages in his book suggest it, Craig does not explicitly endorse high standards invariantism. Rather, he allows that the meaning of a knowledge attribution might be able to undergo variation in actual use. To explain the possibility of this occurrence, he refers us to Unger's account of semantic relativity.[194]

According to Unger, what is communicated by a knowledge attribution varies according to context. That communication might be affected by a variation in the semantic content of the word "knowledge" or it might be affected by a variation in what is pragmatically implied by the speaker. That is, "knowledge" might have an invariantist semantics and contextualist pragmatics, or it might have contextualist semantics and invariantist pragmatics. Unger argues, and Craig accepts, that the data will never be able to conclusively demonstrate which allotment of contextual variation is correct. Therefore, the proper position is semantic relativity: perhaps a knowledge attribution operates contextualist semantics, and perhaps not.

I explained in 4.b that a functional synthesis will be satisfied by either pragmatic or semantic contextualism. Therefore, I refrain from arguing for one or the other. Here I introduce pragmatic contextualism only in order to demonstrate that it is no less vulnerable to the authority and intelligibility objections than semantic contextualism.

ii. Loose use

Wayne A. Davis is a pragmatic contextualist. He claims that the concept of knowledge has one fixed semantic meaning, but that what is pragmatically implied by a knowledge attribution may vary contextually thanks to the phenomenon of "loose use."[195]

> ...when we use a term loosely in a sentence, we imply that the proposition it expresses is a good enough approximation to the truth so that the difference does not matter for current purposes.[196]

Davis offers the following example as a paradigmatic case of loose use.

[194] (Unger, 1984)
[195] (Davis, 2007, p. 408)
[196] Ibid., p. 407.

> The Coffee Case
> A. When the scoop comes up empty in the coffee jar, I yell to my wife, ‘‘The coffee is all gone.’’ B. When my son comes down for breakfast a few minutes later, he announces that he needs a few coffee grounds for his science project, and then asks ‘‘Is the coffee really all gone?’’ I say with no embarrassment, ‘‘No, there may be enough for you.’’[197]

In context A, Davis uses “all gone” loosely, only implying that the coffee is close enough to “all gone” to satisfy the purpose of brewing another cup. In context B, he has to use the term “all gone” more strictly. Strictly speaking, “all gone” means “not one little bit left,” and so, strictly speaking, the coffee is not all gone and his statement to his wife is a falsehood. Yet he can speak loosely to his wife without actually communicating a falsehood, because he is sure that she will understand him loosely as well. She understands the proposition he *meant* rather than the proposition he actually *said.* When, with his son, the context changes so that he realizes that he will be understood strictly, according to what he actually says, then he must speak strictly to avoid communicating a falsehood.

According to Davis, our use of the word ‘knowledge’ follows the same principle. Sometimes we may use the word knowledge just because it is close enough to real knowledge to count in the given context. We say “knowledge,” but we pragmatically implicate something less. The feature that varies according to context is how much divergence from the true semantic meaning of knowledge is acceptable.[198] If the situation has low stakes, we do not object if the speaker uses the word very loosely; we understand that by saying “S knows,” he actually only means that S’s belief is somewhat justified. If the situation has high stakes, we expect the speaker to use the word with precision; that is, he should only say “S knows” if he means that S has enough justification to really know, strictly speaking.

iii. Pragmatic contextualism also loses authority and intelligibility

Whether the contextualist sensitivity is pragmatic or semantic, it produces contextually varying effects in its hearers. That is, whether or not the semantic meaning of the word ‘knowledge’ changes, what is actually communicated and understood does change. Therefore it is unclear how attributions could operate intersubjectively, remaining authoritative and intelligible in cases of

[197] Ibid., p. 406.
[198] Ibid., p. 408.

recommendation. If someone recommended an informant to me, or if I overheard an attribution, I would need a good deal of background information to be able to decode the meaning, whether that information changed the way I understood the semantics or the pragmatics of the speech act. Any semantic invariantists who argue that contextual variation is conveyed by pragmatic implications must explain how their account nonetheless retains intersubjective authority and intelligibility.

Davis' loose-use account therefore runs afoul of the disquotation objection. Rather than decoding the speaker's *semantic meaning*, an inquirer is left with the task of decoding the speaker's *pragmatic implication*, but there is no reason to expect that to be any easier of a task. We still cannot understand how people can report knowledge attributions, yet avoid a widespread rash of misunderstandings.

Davis does not realize that his account is vulnerable to the disquotation objection. In fact, he actually offers the disquotation objection against semantic contextualism, claiming that the following statement is valid in all contexts: "S speaks truly when saying I know p. Therefore, S knows p."[199] Such a statement can only be reliably valid if knowledge attributions are disquotable. Otherwise, S's true statement might be untrue relative to the speaker; compare "Janice speaks truly when saying Jack is nearby. Therefore, Jack is nearby." In many contexts, namely all contexts in which the speaker is at a different location than Janice, that statement is invalid.

Davis assumes that pragmatic contextualism allows the statement to be valid in all contexts. Thus pragmatic contextualism scores a point over semantic contextualism, for it can account for the disquotability of knowledge attributions. In fact, pragmatic contextualism is unable to ensure that the statement will always be valid. Davis claims that people use the word knowledge while pragmatically communicating some less demanding state with great regularity.[200] Therefore, it will be the case that a given knowledge attribution is very likely to be, strictly speaking, false. In any of these common situations, the speaker of the above statement may be pragmatically implying something less than 'know' when she says, "S speaks truly when saying I know p." Since knowledge is so rare, it would be an entirely

[199] Ibid., p. 400.
[200] Ibid., p. 428.

acceptable and rational conversational move to use that sentence to pragmatically *imply,* "S speaks close enough to the truth when saying I know p."[201]

Now we should ask, *what's to stop the speaker's pragmatic context from changing mid-statement?* Perhaps with the second sentence, she drops the pragmatic implications and speaks entirely literally. Therefore, she could acceptably say, "S speaks truly when saying I know p. However, S does not know p." Thanks to the pragmatic implications that modify what is communicated by the first sentence, the whole statement is still valid.

Davis attempts to block this consequence by stipulating that the pragmatic implications of a word cannot change within the same statement.[202] However, he does not offer any support for this stipulation. Furthermore, there is no reason to assume that the exact same move would not be open to the semantic contextualist who could also simply stipulate that a context may not change within the same statement. I will suggest in section 6.a that the context of a sentence indeed may not change mid-sentence, at least not without a lot of guiding clues to alert the listener. However, this fact does not, in itself, offer any guidance as to whether contextualist sensitivity is pragmatic or semantic.

As stressed in 4.b, I do not mean to argue that the contextualist variation of knowledge attributions must be semantic. By criticizing Davis' loose use account, I mean rather to show that one may not avoid the authority and intelligibility objections simply by shunting contextual variation into pragmatics. If any contextual sensitivity is possible, pragmatic or semantic, we must understand how we manage to avoid the loss of intersubjective authority and intelligibility.

Verdict

Contextualist sensitivity, whether semantic or pragmatic, allows knowledge attributions to function appropriately to their practical environment only by obscuring the possibility of intelligibility and authority.

[201] Note that this possibility does not imply that the semantic meaning of 'truth' or 'truly' may vary contextually. Rather, it implies that people may in speaking use 'true' loosely.
[202] (Davis, 2007, pp. 409 - 410)

e. Convention contextualism

i. A middle way

All the thinkers presented above seem to assume that we must choose between two extremes. Either we index the concept to one individual at one time, thereby allowing for constant and rapid fluctuation and the consequent loss of intelligibility and authority, or else we fix the concept to all of humanity for all times, thereby allowing for absolute rigidity and the consequent loss of sensitivity. However, we needn't restrict ourselves to these two poles. In what follows, I argue that, while the concept of knowledge is less stable than invariantism would have it, it is fixed to practical environments that are more stable than one individual's idiosyncratic preoccupations. The practical environments to which the concept of knowledge is sensitive must be the product of social conventions. By indexing the concept of knowledge to publically recognized contexts, we may retain sensitivity without subjecting the implied degree of epistemic authority to the fleeting whims of any individual subject. Furthermore, given the public nature of conventional contexts, listeners nearly always have enough background information to be able to understand the attribution, which accounts for the concept's general intersubjective intelligibility.

This promising approach has been suggested, or at least hinted at, in three recent essays. Greco, Willaschek and Beyer have all lent support to the notion that knowledge attributions are modified by social, rather than individual, practical environments. By taking the best part of each of their approaches, I present an approach that can synthesize all three of the desiderata.

ii. Constraints on interest-dependence: Greco

John Greco arrives at contextualism by means of his position that, to count as knowledge, a belief must be saliently caused by the believer's intellectual ability.[203] By way of accepting contextualism about salient causal relationships, he accepts contextualism about knowledge attributions. Just I as have done in the last chapter, he remarks that Craig's work suggests that the relevant context is determined by the

[203] Greco 2009.

demands of practical reasoning and uses Hawthorne's notion of the practical environment to explore the consequences.[204]

In "What's Wrong with Contextualism" Greco acknowledges that contextualism might seem to threaten the objectivity of knowledge by making it interest-dependent. If the truth value of a knowledge attribution were interest-dependent, we might find that "odd or disturbing (or ugly, or we just wouldn't like it)."[205] To mollify the distastefulness of interest-dependence, Greco introduces two constraints on the kinds of interests that may be allowed to affect the truth value of a knowledge attribution.[206] First, the interests must be the interests of a group, rather than an individual. Second, the interests must be actual rather than perceived.[207]

With Greco's constraints, knowledge attributions are sensitive to actual group interests as opposed to perceived individual interests. Such an approach obviously holds great promise for my project of synthesizing the conditions of sensitivity, authority and intelligibility. Under these constraints, the concept of knowledge remains sensitive to practical environments; the difference is that the practical environments must be *publically acknowledged*. If the concept of knowledge is sensitive only to public practical environments, then it seems possible that the intelligibility and authority conditions could be satisfied while retaining a high level of flexibility. A restriction to public practical environments might assure that the relevant practical environment would always be apparent, no matter who was making the attribution. This apparentness would demand detectable clues that transcended individual speakers. What would such clues look like?

iii. Epistemic practice contextualism: Willaschek

Greco's constraint is promising, but we need a more detailed picture of what constitutes a public practical environment. In "Contextualism about knowledge and justification by default,"[208] Marcus Willaschek offers just that, outlining a new model

[204] Greco 2008, p. 433.
[205] Ibid., p. 430.
[206] Ibid., p. 431.
[207] See also Dretske 1981, p. 377. "...the difference between a relevant and an irrelevant alternative resides, not in what we happen to *regard* as a real possibility (whether reasonably or not), but in the kind of possibilities that actually exist in the objective situation."
[208] Willaschek 2007.

of contextualism which he calls epistemic practice contextualism. To set the stage for the new model, Willaschek draws on the work of Michael Williams.

Williams does not support contextualism by using semantic data of the use of the word 'knowledge'; indeed, his points should hold for cases where the knowledge-claim is only implicit, and does not involve the actual word at all.[209] He argues that skeptics fail to recognize that, at any given point in time, some beliefs must have default status; they need no argumentation, but may simply be accepted. He describes default status using two forms of justification, both of which are necessary for knowledge: personal and evidential. Personal justification is a question of whether the subject's belief is *responsibly formed*. Absent any challenges to the contrary, Willaschek argues that subjects' beliefs count as responsible by default, "Personal justification is ... like innocence in a court of law: presumptive but in need of defence in the face of contrary evidence."[210] Evidential justification, on the other hand, is an externalist consideration; a belief can be responsibly formed, and thus personally justified, but nonetheless evidentially unjustified thanks to circumstances beyond the subject's control. Here is an example: buying a genuine type-E vintage Jaguar is much cheaper than building a replica. Ivan knows this and thus, his claim that the type-E Jaguar before him is a genuine vintage car is responsibly formed. However, unbeknownst to Ivan, some wealthy enthusiast has been making replicas of type-E vintage Jaguars.[211] This fact makes Ivan's claim evidentially unjustified, despite the fact that it was responsibly formed. How many cars must the enthusiast have made to destroy the evidential justification for Ivan's belief? According to Williams, whether a belief counts as evidentially justified depends on the specifics of a given context. The actual demands are often hard to pin down, because "there is considerable indeterminacy about the objective adequacy of grounds, resulting from the fluidity of contextual boundaries."[212]

As I have been stressing, this fluidity of contextual boundaries has spurred the concern that, on a contextualist account of justification, the authority and intelligibility of intersubjective knowledge claims is inexplicable. However, Williams does not address this problem. Neither does he investigate in greater detail

[209] Williams M., 2001, pp. 26 – 27.
[210] Ibid., p. 25.
[211] Ibid., p. 163.
[212] Ibid.

how various default statuses can differ from one another. He does outline five features that can influence a claim's epistemic status, but they do not bring us much farther than the three factors of practical environments sketched in 4.b.[213] Given his concern with skepticism, Williams' goal is not to differentiate between different sorts of contexts, but rather to demonstrate how any default status may remain immune to skeptical doubts, which have no context at all. However, this leaves us without an entirely satisfying reply to the authority and intelligibility objections. Granted that knowledge claims change their character according a mesh of various factors, how can we act on and understand the claims?

Willaschek develops Williams' account by arguing that the default status of a belief is relative to the established social epistemic practice of which it is a part. Some epistemic practices are highly regimented; they offer explicit rules according to which practitioners must "acquire and attribute" knowledge.[214] There are also less self-reflective practices, such as "crafts, commerce, and sports."[215] While these practices mostly share the same "all-purpose set of epistemic standards that governs commonsense attributions of knowledge,"[216] they also employ some standards peculiar just to their own practitioners. Willaschek offers some examples:

> The epistemic standards employed in different practices overlap, but there are also important differences. In the empirical sciences, for instance, knowledge is tied to the possibility of empirical confirmation; in mathematics and related formal disciplines, knowledge requires proof; in the law, knowledge from testimony is restricted by certain formal procedures such as taking an oath; in various crafts, practitioners can tell things apart simply by looking or touching, while laypersons can do so only by indirect methods; etc. etc.[217]

When analyzing a given institution to work out what sort of standard it places on knowledge, we should recall our basic hypothesis and the explication it produced in KAP. Imagine an inquirer, Ella, wondering whether Louis qualifies as having epistemic authority. How does the relevant epistemic practice affect what conditions

[213] In addition to end of inquiry and stakes, he mentions two kinds of relevant alternatives: situational and conversational, and the presuppositions of semantic intelligibility. (Ibid., pp. 159 – 162) The *situational* factor brings us closest to what we need, as it entails an objective external factor against which a subject can make mistakes, but Williams does not clarify in more detail what governs our recognition of this factor and the ways in which publically recognized conventions can affect it.
[214] Willaschek 2007, p. 263. The suggestion that there is contextual variation not only in knowledge attributions, but in methods of knowledge acquisition is very interesting and plausible. Unfortunately, I've no space to pursue it further here.
[215] Ibid., p. 263.
[216] Ibid., p. 263.
[217] Ibid., pp. 263 – 264.

Louis must fill in order for Ella to trust him and accordingly use the proposition he asserts as a premise in rational practical reasoning? In the passage above, Willaschek makes helpful concrete suggestions. If Ella is working in the empirical sciences, Louis may count as knowing if his assertions are backed by empirical replicable data. If Ella is asking a mathematical question, Louis may count as knowing if his assertions are backed by deductive proof. If Ella is in a courtroom, Louis may count as knowing if his information is based on the testimony of a witness under oath. The list of examples could be extended and specified at length.[218] We should notice that each context bears a detectable clue, such as the attributor's *location* (a courtroom) and the *content* of her question (mathematical). These detectable clues promise to solve the intelligibility objection. Knowledge claims are intelligible because they always hold more pertinent information than a bare statement, "S knows that *p*." By picking out a number of various aspects of and surrounding a knowledge claim, a listener gets an idea of the context to which it is sensitive.

Epistemic practice contextualism manages to satisfy the authority condition as well. For any given context, there is only *one* answer as to whether or not a subject knows. In the courtroom, there is no danger that some attributor will make a knowledge claim using tavern standards; if she did, she has simply made a mistake, a false attribution. Thus, true knowledge claims will always be authoritative for their context. Since public clues make the context always apparent, there is no danger that a claim made in one context will be repeated in another context. People can disquote knowledge attributions only when enough of the surrounding clues remain to ensure that the same standards are in place.

Thus, a knowledge claim can be both context-sensitive and absolute. As Willaschek explains, when a knowledge claim is authoritative for a given context, it has default status. Each epistemic practice sets its own criteria for the amount and kind of justification required for given beliefs to have default status: in a courtroom,

[218] Here I would like to add only one further example, the context of early education. As Alvin Goldman (1999, pp. 363 - 366) points out, in the first levels of education, students cannot be overly critical towards their teachers' assertions. Otherwise, they will not be able to acquire the skills they need for critical thinking in the first place. Therefore, in the context of early education, in which students are inquirers and teachers are informants, the standards for establishing trustworthiness must be laxer than in many other epistemic practices. We might then expect or demand a counterbalancing *stringency* of standards when it comes to hiring the teachers, to ensure that their easy authority does not end up detrimental to the trusting students.

certain assertions may have default status if the speaker would swear to it on oath. In a laboratory, certain assertions may have default status if the speaker can point to reproducible data collected by empirical experiments. According to Willaschek, if a belief is true and has default status, then it is known. If beliefs do not have default status, then they must be explicitly argued for. The point at which a belief *without* default status may be considered established through explicit argumentation is also relative to the standards of the relevant epistemic practice.

As shown in Williams' example of the sports cars, the codified standards of one practice underdetermine the relevant criteria. Ivan's belief does not count as knowledge because of a very specific fact: one enthusiast has been making replicas. Willaschek contrasts such circumstances with the codified standards, calling the former the "facts of the matter."[219] The facts of the matter affect which relevant alternatives a subject would have to rule out for a speaker's knowledge attribution to be true. The standards of the practice can determine in general what *sorts* of possible alternatives need to be ruled out in order for a knowledge attribution to be true, but the facts of the matter are needed in order to judge whether certain *specific* alternatives need to be ruled out. For example, Willaschek claims that ordinary epistemic practice has the following standard: "rule out any error-possibilities for which there's a reason to think that they may in fact obtain."[220] Whether or not there is such a reason is then a fact of the matter. Before there were any counterfeit Rolex watches getting hawked, reading the brand name on the watch face sufficed to grant a belief that the watch was a Rolex default status. After the possibility of counterfeit watches has become not just a logical possibility but a living threat, then the same evidence no longer suffices. In order to judge whether it is a relevant possibility that the watch is a counterfeit, we need to consider the fact of the matter: whether or not there are counterfeit Rolex watches in general distribution.[221]

According to Willaschek, certain paradigmatic kinds of belief may enjoy *universal* default status. If a subject can claim that his belief is based on memory, perception or testimony, then he need not offer any further evidence to convince his listeners to accept it as knowledge. That is, ordinary epistemic practice awards

[219] Willaschek 2007, p. 264.
[220] Ibid., p. 265.
[221] Ibid., p. 264.

default status to beliefs acquired by memory, perception and testimony.[222] He points out that even these universally acceptable beliefs can lose their default status if a fact of the matter indicates that they are no longer reliable sources. That I remember doing something means that I know I did it, unless it so happens that I have dementia.

A default belief only loses its status if there is a reasonable challenge to its reliability. What counts as a reasonable challenge also varies according to the epistemic practice and the facts of the matter.[223] Thus, it will be relative to the context at hand how many counterfeit watches must be in circulation before reading the brand name on its face no longer awards the belief 'this watch is a Rolex' default status.

In sum, four points can change according to context:

> (a) what kinds of beliefs enjoy default status; (b) what kinds of error-possibilities are relevant challenges; (c) what counts as answering a challenge (as ruling out an error-possibility); and (d) what counts as establishing a claim that does not enjoy default status.[224]

These four questions are answered in principle by the epistemic practice's standards, which in conjunction with the facts of the matter decide the status of specific beliefs. With Willaschek's epistemic practice contextualism, we now have a much clearer blueprint of how to develop Greco's proposal and index the meaning of knowledge to the actual interests of groups rather than the perceived interests of individuals in a way that allows for public authority and intelligibility.

However, Willaschek makes a further argument which would drastically impact the sensitivity of knowledge attributions. He argues that for any given claim, there is either an expert context or there is not. If there is a strict expert context, then the knowledge claim must be evaluated according to its standards; if not, 'everyday' standards apply.[225] On the strength of this dichotomy, Willaschek denies that a knowledge claim could ever be true in the mouth of one attributor, while false in the mouth of another. If Renata claims to know that it will rain by looking at the clouds, and Maria the meteorologist denies that Renata knows, than only Maria's claim can

[222] Ibid., pp. 257 – 258.
[223] See also Dretske 1981, pp. 374 – 375. Dretske argues that the relevance of alternatives is sensitive to context and that the mere suggestion that a source of information *could* be wrong always counts as an irrelevant alternative.
[224] Willaschek 2007, p. 264.
[225] Ibid., p. 267.

be correct, for only Maria is using the expert standards of the relevant epistemic practice. Thus, epistemic practice contextualism does not allow for the kind of sensitivity that my synthesis demands. Since stakes, for example, cannot affect whether or not there is an expert practice, Willaschek implies that stakes cannot affect the truth of a knowledge attribution. Furthermore, with this clause Willaschek contradicts Williams as well, who holds that factors such as stakes are indeed relevant to whether a subject counts as knowing or not.

iv. Conventions rather than practices: Beyer

In order to synthesize the sensitivity, authority, and intelligibility conditions into one coherent concept, we need to think of public epistemic contexts, but we must also allow for these public contexts to change according to factors such as stakes. In "Contextualism and the background of (philosophical) justification," [226] Christian Beyer presents a key concept which can enable this synthesis: epistemic *conventions*. Conventions are, of course, public, but unlike Willaschek's practices, they need not be rigidly defined. Conventions can be based on subtle social cues. Beyer introduces the apt phrase 'convention contextualism' and defines it as follows:

> *S* knows that *p* at time *t* for a speaker belonging to a given linguistic community *l*, in a context of assessment *c*, iff (1) *S* believes at *t* that *p*, (2) it is true that *p*, and (3) *S* meets, at *t*, the criteria of justification that, according to the conventions valid in *l*, apply in *c*.[227]

Beyer does not offer a clear explanation of what constitutes a convention. He claims that conventions are generally shared presumptions as to which alternatives are epistemically possible,[228] but does not pursue the notion further, referring us instead to the work of two other philosophers: Husserl's phenomenological analysis of the *Lebenswelt* (life-world) and Wittgenstein's discussion in *On Certainty* of a subject's background or *Weltbild* (picture of the world.)

Despite the vagueness of what constitutes a convention, I expect that this notion will allow us to helpfully modify Willaschek's model of epistemic practice contextualism, using its insights to meet the authority and intelligibility objections, without overly narrowing the scope of acceptable contexts. The synthesized concept

[226] Beyer, 2007.
[227] Ibid., p. 295.
[228] Ibid., p. 297 – 298.

can demonstrate sensitivity to subtler and vaguer, but nonetheless public conventions. We may then accept a change in stakes as affecting a change in the meaning of knowledge attributions, for people conventionally do adjust their epistemic standards to the stakes of their situation. Therefore the statement, "He knows that this mushroom isn't poisonous" might conventionally express a higher degree of reliability than, "He knows that this mushroom is one dollar cheaper than the other," for the former implies a recognition of life or death stakes and we humans tend to care a lot more about staying alive than about saving a dollar.

Stakes are not the only factor unaddressed by epistemic practice contextualism that can nonetheless affect what will conventionally be understood by knowledge attributions. Take Bach's eavesdropping objection.[229] If you overhear someone say, "Nixon knew that Liddy was planning the Watergate break-in," do you understand him because you grasp which *practice* governs the standards of his attribution? While politics may be described as a practice, I do not believe that this is the best way to approach the intelligibility of the assertion. Rather, you gather that he is talking about a high rank political scandal. Given what people understand about high rank political scandals, the statement conventionally implies both very high stakes and a very low level of required justification. That is, we do not expect that the speaker means that Nixon was able to prove beyond any shadow of a doubt that Liddy was planning the Watergate break-in. Given the stakes and the nature of the situation, Nixon may be publically attributed with that knowledge as long as he had *any* substantial evidence at his disposal. On the other hand, a statement such as "President Obama knows that they harbor nuclear weapons," while implying equally high stakes, would imply that Obama has a *huge* amount of substantial evidence, otherwise the public would not be willing to attribute him with knowledge of such a thing.

I do not expect these examples to be immediately convincing. I only mean to demonstrate how one could begin to apply convention contextualism in order to describe these situations. As I will discuss at more length in the next chapter, these musings must be supplemented by linguistic and sociological research as well as normative arbitration.

[229] Bach 2005, p. 61.

v. Verdict for convention contextualism

Convention contextualism provides us with the tools to synthesize all three conditions into one coherent concept. It limits the amount of sensitivity, but not enough to disable the concept from playing the hypothesized function in practical reasoning. While epistemic conventions do not respond to the practical environments of isolated individuals, they do adjust to publically recognized practical environments. Therefore, people may use the concept of knowledge in a way appropriate to their situation, as long as their situation fits some generally acknowledged mold.

According to convention contextualism, the meaning of a knowledge attribution cannot change according to entirely personal factors, such as whether or not the speaker happens to be thinking of a certain counterpossibility when he makes it. Rather, the meaning will be indexed to generally acknowledged standards. As indicated above in my discussion of Bach's eavesdropping example, the *content* of knowledge attributions will often indicate the conventional context according to which they ought to be understood. When a person makes her attribution according to conventional standards, listeners can thus understand it according to those same standards, even without any further background knowledge. Therefore, there is little danger that the listener will find herself using the 'known' premise as a reason for action in a situation in which her raised standards make that no longer a rational choice. Within one fixed conventional context, there is one publically recognized standard for a belief to have default status. Beliefs that achieve that default status, or that are maverick but explicitly grounded, may be considered as absolutely authoritative within that context. Since anyone who hears the knowledge attribution should understand the conventional context it implies, it is possible for the attribution to imply a generally acknowledged status, honor, and normative demand as well.

f. The final synthesis: KAP–CC

The consideration of competing syntheses of knowledge have shown that, to accommodate the two objections and fully account for the social demands highlighted by my hypothesis, we should modify KAP by indexing the reliability

condition to S_1 - S_n's conventional practical environment. The resulting formulation looks as follows:

Knowledge characterized by its Authoritative role in Practical reasoning - sensitive to Conventional Contexts (KAP–CC)

S_1 knows relative to S_1 - S_n iff *p,* S_1 believes that *p,* and S_1 has a property X which correlates with being right about *p* reliably enough to count as authoritative by the standards of S_1 - S_n's conventional practical environment.

KAP–CC is the final result of my conceptual synthesis. In the following chapter, I shall compare it to commonly held intuitions about knowledge and test for discrepancies.

6. Comparison: Is knowledge equivalent to KAP – CC?

In the following chapter, I compare the result of my conceptual synthesis to intuitions about the concept of knowledge. I consider three possible discrepancies and argue against each of them. In so doing, I also manage to show why KAP–CC is superior to Hawthorne and Stanley's subject-sensitive invariant conception of knowledge.

In conclusion I discuss the vagueness about what constitutes a conventional practical environment. While we need not strive to eliminate all vagueness, I suggest that here there is room for empirical research and normative modifications as well as factual and conjectural genealogies.

a. The disagreement objection

i. The alleged contextualist interpretation of disagreement

The disagreement objection points out that any contextualist conception of knowledge does a poor job of accounting for intuitions generated by disagreements over whether or not someone knows.[230] Since it is usually leveled at those contextualist accounts motivated by skeptical arguments, the objection is usually based on examples of a skeptic arguing with a non-skeptic. So, for example, Mark Richard takes the example of Moore arguing with a skeptic about whether he knows that he has hands.[231] Moore says he knows it; the skeptic says he does not. Contextualist accounts supposedly must interpret this situation as follows:

> The two speakers do not really contradict one another at all. The whole dispute is based on a misunderstanding! If the two speakers only clarified what each of them meant by the word know, they would realize that they are actually making two different claims. Knowing relative to the skeptic is simply a different property than knowing relative to Moore. Therefore, they are simply talking past one another, and both may very well be right, for Moore might know relative to himself and not know relative to the skeptic.

Since KAP–CC is also a contextualist account of knowledge, it presumably would be forced to give just such a gloss on the disagreement. Yet, this interpretation does not sit well with our intuitive response to the disagreement. As Richard puts it,

[230] See Bach 2005, p. 53; DeRose 2009, pp. 129 – 131; Hawthorne 2004, p. 191; Richard 2004; Rysiew 2001, p. 485.
[231] Richard 2004, pp. 215 – 216.

> One feels that something is awry. One wants to say that when the skeptic and Moore argue with each other, they disagree about whether Moore knows he has hands. One wants, for that matter, to say that when the skeptic utters 'Moore doesn't know he has hands' he says that Moore doesn't know that he has hands; analagously, for Moore and 'Moore knows that he has hands'. But it's not clear we can say any of this, if we are contextualists.[232]

The two parties believe that they are actually disagreeing and that only one of them can be right. As long as we have not already converted to contextualism, our intuitions concur. KAP–CC thus poses a discrepancy with widely held intuitions about the way the concept of knowledge is actually used.

According to DeRose, no advocate of contextualism would actually endorse such an interpretation of disagreement.[233] There are various ways to offer a contextualist interpretation of disagreement without arguing that both speakers can be speaking truly. I will discuss here only that which convention contextualism offers.

ii. Practical environments must be shared by both speakers

My synthesis has not been driven by the desire to interpret skeptical arguments. What sort of practical environment could the skeptic have in mind? It is thus not immediately clear how the KAP–CC synthesis would guide us in describing Moore's and the skeptic's disagreement. Of course, its weakness in dealing with knowledge claims that do not relate to practical reasoning is itself another objection to KAP–CC which I return to in section d below.

To respond to the disagreement objection, I must consider disagreements that could arise out of two different practical environments. What sort of disagreement is possible in such cases? As demonstrated in the example from Chapter Four, "John was at work," sometimes what looks like the same knowledge attribution may actually be understood as two different propositions with two different truth values. This apparent contradiction arises because the propositions are asserted in two different practical environments with two different standards for epistemic authority. If Thelma and the police were to disagree about whether or not Louise knows that John was at work, they would indeed merely be talking past one another. However, notice that in the example Thelma and the police are not actually speaking to one

[232] Ibid., p. 216.
[233] DeRose 2009, p. 131.

another. Their 'disagreement' is only an artificial possibility. That is, if some *neutral observer* considers the two statements next to each other, it may seem as if the speakers would disagree, were they to start a conversation. But in fact, were the two speakers to confront one another and actually enter into a dispute, then they would not be able to avoid being in the same practical environment. If the police began to question Thelma as a potential witness to John's whereabouts the day of the crime, Thelma should adapt to the conventions that govern criminal investigations. If Thelma were to ask the police for their advice as to whether or not she should pay up to Louise, the police should adapt to the conventions that govern petty bets.

This response to the disagreement objection relies on the following important condition on the public practical environment: it must be shared by both speakers. In a normal conversation about whether or not someone knows that *p*, both speakers adapt to common conventions as to what knowledge means. Since two people need to speak to one another in order to actually disagree, KAP–CC will never offer the alleged interpretation caricatured above.

A natural way to phrase this condition is to say that conversations cannot bridge contexts. However, it would be false to claim that one cannot be in two conventional contexts at the same time. DeRose offers the example of Jane, the walking talker.[234] Talking to her friend, she claims to know that she will return to the college next year, despite the fact that she is that very moment walking to go sign up for a year of life insurance and thereby acting on the premise that she does *not* know any such thing, for she cannot rule out the possibility that she will die before the year is out. The woman uses the concept of knowledge in two different ways within one conversation. Relative to her conversation with her friend, she *does* know that she will return to the college next year. Relative to her decision to buy life insurance, she does *not* know that she will return to the college next year.

Jason Stanley argues that such shifts within one conversation are not possible.[235] However, the linguistic data on which he bases his case uses only one proposition, with no extra details that imply any *reason* why the relevant context would shift. Conventions, though, rely on just such clues. When we look at a more thoroughly described example, such as DeRose's walking talker, we may recognize the

[234] DeRose 2009, pp. 270 – 271.
[235] Stanley 2004, p. 23.

difference between the two contexts which have generated the two apparently conflicting uses. Chatting about plans obviously entails a different end of inquiry than buying life insurance. Only by being used in two *recognizably* different *conventional* contexts may one get away with bridging a context within one conversation. Thus, these situations should not engender disagreement at all. If the woman's friend objects that the woman does not really know whether she will return to the college, the friend is either mistaken, or attempting to switch contexts. In such a case, we can encounter genuine disagreement.

iii. Genuine disagreement

There can be genuine disagreement within one context. That is, there are situations in which two subjects disagree despite being in the same practical environment, and thereby meaning the same thing by 'knowledge.' If Molly and Rachel are in the same public practical environment, and if Molly believes that "Edgar knows that *p*" while Rachel believes that "Edgar doesn't know that *p*," then only one of them can be right. Perhaps Rachel denies knowledge of Edgar because she overestimates the relevance of a certain alternative. Since her community would deem that alternative irrelevant, then her attribution is false. By leaving open the possibility for such mistakes, convention contextualism thus differs from other forms of contextualism. According to salience-sensitive contextualism, for example, if one individual subject deems an alternative relevant, then so it is; it would thus be impossible for Rachel to mistakenly award relevance to an irrelevant alternative. Thanks to Greco's constraint to *actual* rather than perceived interests,[236] such errors are possible according to convention contextualism and so genuine disagreements may arise.

In genuine disagreements, one of the speakers must have made an error. Taking my cue from Willaschek's epistemic practice contextualism,[237] I suggest that according to convention contextualism, two kinds of errors are possible: an error as to which standards govern the environment at hand and an error as to the facts of the matter.

[236] Greco 2008, p. 431.
[237] Willaschek 2007.

Suppose that, in fact, John was not at work. In such a case, both Thelma and Louise are wrong in the petty bet scenario, thanks to an unexpected fact of the matter.[238]

Now, leaving it an open question for the moment whether or not John was, in fact, at work, suppose that Thelma does not accept Louise's epistemic authority and thereby refuses to give up the five dollars. The two begin to fight about whether or not Louise's evidence is sufficient to resolve the bet or not. "You're being ridiculous; he was obviously at work," argues Louise. "If you didn't actually see him, you can't be sure," argues Thelma. "Come on, we both know he was there!" says Louise. "No, maybe he just forgot his hat yesterday! Maybe Fred was confused, or you misheard him! You can't be sure he was there," says Thelma.

The two of them are in fact operating different standards for epistemic authority. However, they are also genuinely disagreeing: they disagree over which standards should be taken as appropriate for the context at hand, i.e. over which conventions govern their petty bet. With this interpretation, one can be a contextualist and accommodate disagreement based intuitions. Simply point out that the disagreement is not over the truth of the assertion, but rather over the meaning of one of the verbs it contains.[239] Thelma and Louise do not so much disagree about the answer to the question "Does Louise know or not?" as about the answer to the question "In our petty bet context, is circumstantial evidence enough to count as knowing or not?"

So much for the intuition that they are in disagreement. What about the intuition that only one of them can be right? Who is right, Thelma or Louise? In responding to this intuition, we must acknowledge the *vagueness* of the conventions pertinent to this example. Here, as in most other conventional practical environments, there is no well-defined code of standards to which any knowledge attribution must adhere. Rather, the conventions are flexible and vague. Therefore, it is true that in this example of disagreement, there is no clear answer as to which of them is right. While gambling institutions have established standards, a friendly personal bet is a 'gentleman's agreement,' made on the implicit assumption that both parties recognize the same conventional standards of epistemic authority. If it turns out that

[238] Was it therefore inappropriate for Thelma to give Louise the five dollars? I discuss my position on behavior evaluation in section b below.

[239] See Richard 2004, p. 217. Richard argues that this is the intuitive way to understand disagreement, presenting the argument as an objection to (individualistic) contextualism.

they both have different expectations, then intelligibility has failed. In such cases the disagreement is genuine, but not clearly resolvable.

Thus, irresolvable disagreements in which neither party is clearly right or wrong arise only in those cases in which knowledge attributions are neither intelligible nor authoritative. If such cases arose often, then both the functional strength and the intuition accommodation of convention contextualism would be threatened. However, such cases do not arise very often. When they do, I presume that two responses will be the most common, both of which contribute towards reducing the prevalence of such cases. First, the parties can try to convince each other as to what the standards should be. Second, the parties can explicitly state the standards in the future. When vague conventions lead to conflict, the natural and appropriate response is to clarify the conventions.

b. The trans-contextual behavior evaluation objection

i. The alleged contextualist interpretation of behavior evaluation

Hawthorne argues that no contextualist account can accommodate the function of the concept of knowledge in behavior evaluation scenarios.[240] His argument proceeds by showing that the following three claims cannot all be true:

(1) Whether or not S_1 may use *p* as a premise in rational practical reasoning does not depend on S_2's practical environment.
(2) If S_1 may not use *p* as a premise in rational practical reasoning, then S_1 does not know that *p*.
(3) Whether or not it is true that 'S_1 does not know that *p'* depends on S_2's practical environment.

One cannot simultaneously affirm all three claims. To understand why, imagine that the police are talking about Thelma. They recognize that, in *her* petty bet context, (1) Thelma was acting rationally to trust Louise and pay the five dollars. However, according to *their* criminal investigation context, they agree that (3) "Thelma does not know whether or not John was at work." Thus, they must be able to say (not-2), "Thelma may act as if John was at work; nonetheless, she does not know that John

[240] (Hawthorne, 2004, p. 90)

was at work." In this case, the following two statements are true, with S_1 as Thelma and S_2 as the police:

(1) According to S_1's practical environment, S_1 may use *p* as a premise in rational practical reasoning.

(3) According to S_2's practical environment, S_1 does not know that *p*.

Given the truth of (1) and (3), S_2 must be able to make the following statement,

(not-2) "S_1 may use *p* as a premise in rational practical reasoning, but she does not know that *p*."

In this situation, affirming claim (1) and (3) leads to the denial of claim (2). In the police's context, it is true that 'Thelma does not know that *p*.' However, thanks to Thelma's context, it is false that 'Thelma may not use *p* as a reason for action.'

Claim (2) is the basis of the role of the concept of knowledge in practical reasoning, according to both subject-sensitive invariantism and KAP–CC. Neither theory wants to deny it. Claim (3) is the contextualist thesis. Subject-sensitive invariantism saves claims (1) and (2) by denying claim (3).

Hawthorne imagines that contextualists have to deny claim (1), leaving them with the burden of a more far-reaching contextualism than perhaps they bargained for.[241] By supporting contextualism about "know" they are forced to support contextualism about practical rationality. That is, the contextualist would have to argue for something like:

(4) Whether or not it is true that 'S_1 may use *p* as a premise in rational practical reasoning' depends on S_2's practical environment.

Claim (4) supports contextualism about rational practical reasoning. A contextualist could save claim (3) by replacing claim (1) with claim (4). However, claim (1) has a lot of intuitive support; whether or not S_1 may perform a certain action should depend only on his *own* practical environment. My practical environment has no bearing on how you ought to act in your practical environment. The fact that *Ida's* child is allergic to peanuts does not mean that *Estelle* should be more careful about

[241] Hawthorne 2004, p. 91.

what cookies she eats. Furthermore, behavior evaluation does seem to rely on the truth of claim (1). That is, if John's mother can truthfully tell John "You ought to buy life insurance," then it would be odd if John could just as truthfully reply, "That's true for you, but not for me! According to *my* standards, I ought to save my money!" Contextualism about rational practical reasoning does not seem compatible with behavior evaluation. So, if Hawthorne's criticism is correct, then any contextualist account of knowledge, including KAP–CC, would be unable to fully account for the force of behavior evaluation. In that case, the choice between the two theories would come down to choosing which of the two types of scenario is more important: behavior evaluation or testimony evaluation.

ii. The conventions of trans-contextual behavior evaluation

Hawthorne's criticism of contextualism is incorrect. There is a form of contextualism that can account for the function of the concept of knowledge in behavior evaluation, namely, convention contextualism. KAP–CC can describe behavior evaluation without succumbing to contextualism about rational practical reasoning. We must simply recognize that, according to the conventions of fair behavior evaluation, you may not evaluate another's behavior without adapting your standards to the context of the subject under evaluation. If S_2 is evaluating S_1's behavior, that fact affects S_2's own context. Conventions of fairness hold that S_2 may not evaluate S_1's use of *p* as a reason for action by standards other than those pertaining to the environment in which S_1 behaved.[242] That is, conventions can demand that the attributor be sensitive to the subject. Accordingly, when evaluating S_2's behavior, S_1 cannot accurately evaluate S_2's relevant knowledge by standards other than those pertaining to the environment in which S_2 behaved. If the police are judging Thelma's behavior, they should adapt their standards to Thelma's environment. They should set their own high stakes concerns aside and look at things from Thelma's point of view. Otherwise, they cannot give a fair evaluation of her behavior.

Thus, S_2 cannot truthfully say, "S_1 may use *p* as a reason for action, but S_1 does not know that *p*."[243] Convention contextualism can affirm (1), (2) and (3). It avoids

[242] For similar arguments in favor of contextualism's ability to account for behavior evaluation scenarios, see DeRose 2004, p. 349; DeRose 2009, p. 250; and Greco 2008, p. 425.

[243] If S_2 explicitly refers to the divergence in contexts, she can get away with differentiating between S_1's knowledge and her practical reasoning. The police could say, "Given Thelma's situation, it's

those situations in which such their simultaneous affirmation would entail a contradiction by putting a constraint on which practical environment S_2 may be in when evaluating S_1's behavior.

iii. Subject-sensitive invariantism and testimony evaluation

KAP–CC is in fact compatible with the function of the concept of knowledge in behavior evaluation. However, subject-sensitive invariantism is not compatible with the function of the concept of knowledge in *testimony* evaluation. That is, of the two accounts, only KAP–CC is flexible enough to describe the following situation:

> An inquirer's own epistemic position with respect to *p* is not very strong. Thus, she turns to an informant, who has an indicator property which correlates reliably enough with being right about *p* to give him the epistemic authority she needs. He asserts that *p*. She correctly judges that he knows that *p* and believes that *p* on his authority. Accordingly, she uses *p* as a premise in her practical reasoning.

In this situation, the concept of knowledge functions when the inquirer decides that her informant has the epistemic authority she needs. Subject-sensitive invariantism on the other hand can locate a functional role for the knowledge attribution only in the final sentence, once the inquirer herself believes that *p* and uses it as a reason for action. It cannot account for the role that knowledge plays in her evaluation and trust of her informant.

Subject-sensitive invariantism cannot account for the role of knowledge in cases of testimony evaluation. KAP–CC can account for the role of knowledge in behavior evaluation. Thus, according to the criterion of comparison presented in section 4.c.v, KAP–CC is superior to subject-sensitive invariantism.

c. The belief objection

i. Kelp's objection

Craig's work is motivated by situations in which an inquirer needs to decide whether or not to trust an informant. If the inquirer trusts the informant, then the informant's assertion that *p* causes the inquirer to believe that *p*. However, one can easily think

---

appropriate for her to act as if she knew that *p*, but we're conducting a criminal investigation, so we may not assume that she does know that *p*." This would not make claim (2) false, for both sides of the conditional are explicitly contextualized, whereas claim (2) is dequoted.

up examples in which a subject's assertion causes an informant's belief in some deviant way. Kelp presents the following example as trouble for Craig's thesis.

> Henry is in an environment with one real barn and otherwise only fakes. The one real barn is red while the fakes are all green. Henry will believe that a barn is present no matter whether he sees the real red barn or a green fake. However, he is sworn to secrecy about green barns. As a result, in the present circumstances, he will *say* that he is facing a barn only when he is facing a real barn.[244]

Henry, it is generally agreed, does not know that he is facing a real barn. Yet, Kelp argues, he is an objectively good informant as to whether the barn is real.

Now it is easy to point out that this example is bizarre. The tricky part is to explain why it is different than the normal case. As explained above in section 1.d, bizarre counterexamples can be helpful insofar as they illuminate an important feature of the prototypical case that we might not have noticed until we considered a case in which it was absent. When an informant asserts and an inquirer believes, how must the informant's assertion cause the inquirer's belief in order for him to avoid being like Henry?

ii. Craig's distinction between informants and mere sources of information

Craig deals with a similar example by denying that Henry is an informant at all.[245] He argues that anyone believing Henry on the basis of the information given in the example would be using him as a source of information rather than trusting him as an informant. The fact that the inquirer believes that *p* on the basis of the subject's sincere assertion that *p* is still not enough to render the subject an informant. There are many other examples of mere sources of information. Fred's dripping umbrella can give me the true belief that it is raining outside, without making Fred into an informant. [246] Craig also considers how other controversies in epistemological literature may be based on the reluctance to attribute mere sources of information with knowledge.[247]

---

[244] Kelp 2011, p. 8.

[245] Craig responds to the more widely known version from Alvin Goldman, in which there is no talk of green barns and secrecy. However, his response works just as well for Kelp's version.

[246] Craig 1990, p. 35.

[247] He shows how uncertainty over whether someone can know *p* without believing that *p* and whether comparitivism is correct can be explained by the difficulty of distinguishing between an informant and a mere source of information. (Ibid., pp. 37 – 41)

This response would save Craig from Kelp's critique, but is it well justified? If the distinction is to be convincing, we need a functional reason why sources of information should be excluded from the extension of the concept of the good informant.

Craig suggests that an informant is worth distinguishing from a mere source of information, because the former is more convenient for an inquirer. He goes on to mention two features that render the informant more convenient. First, informants are more convenient for the inquirer, because the inquirer does not need any specialized background knowledge in order to interpret the subject's assertion. The difference between a source of information and an informant would be that an inquirer needs to be "specially equipped with background information and inferential techniques not generally available"[248] in order to use a source of information. To understand the informant, on the other hand, the inquirer only needs to understand his language. Second, it is more convenient to have an informant, because he can cooperate with the inquirer, and thereby be more actively helpful.[249]

An inquirer talking to Henry in Kelp's counterexample would need background information about Henry's environment, Henry's awareness of his environment, and Henry's relationship to green barns. We might imagine that this background information is not generally available, and that therefore the inquirer would find Henry less convenient.

iii. The problems with Craig's distinction

The inquirer's need for background information does not adequately capture the difference between informants and mere sources. There are too many situations in which a person does know that *p*, despite the fact that an inquirer would need specialized background information to be able to trust him as to whether *p*. Imagine the following example:

> The inquirer's tribe is at war with a band of marauders, who would all likely lie to her about whether or not a tiger was coming. Only the inquirer has realized that the marauders all have yellow tattoos on their ankles. The inquirer then sees that Fred's ankles are tattoo-free, and accordingly trusts him about the tiger. Without that piece of background information, she should not have trusted Fred.

[248] Ibid., pp. 43 – 44.
[249] Ibid., p. 36.

In this example, Fred knows that the tiger is coming, despite the inquirer's need for specialized background information in order to trust him. There is something else amiss in the Henry example that is not amiss in the Fred example. Furthermore, I suspect that this set of situations, in which a subject knows that *p* despite the fact that people would need specialized background information in order to trust his assertion that *p*, is significantly larger and less freakish than the set of Henry-like situations. Scientists, for example, have specific standards they expect an experiment to respect; they should not trust a colleague's report of an experiment's results without ensuring that her experiment meets those standards.[250] Yet those standards may involve specialized background information. The inquirer's need for specialized background information in order to trust a subject's assertion that *p* does not imply that the subject does not know that *p*. Accepting such a distinction will take care of the discrepancy revealed by the Henry case, but only by producing a whole new batch of discrepancies.

There is a second problem with Craig's distinction: it does not accord with the demands of objectivization. Once we take social and diachronic circumstances into account, we realize that inquirers actually need objectively good informants, rather than subjectively good informants. The objectively good informant, like the subject with knowledge, need be neither accessible nor intelligible to the informant. Even the requirement that the inquirer can understand the informant's language falls away. The concept of the objectively good informant loses in individual convenience what it gains in public usefulness. Yet, as Craig has presented it, sources of information do not count as informants because they are not as convenient to individual inquirers. Why should the distinction between the source of information and the good informant not fall away with the other subjective conditions?

With the loss of the subjective condition (2), the indicator property need not be detectable in the actual world; a possible world in which the inquirer detects it can suffice.[251] Some of those possible worlds will include ones in which the inquirer has just the background information she needs to understand the informant. Therefore, there is no reason why a demand on the inquirer to have background information should exclude an informant from the extension of knowledge. The concept of

[250] See also Keren 2007, p. 371 on the widespread need for background information in normal cases of testimony.
[251] See 3.a.iv.

knowledge, according to Craig's theory, is too objective to make such concessions to convenience.

Furthermore, the subjective condition (4) reads, "Channels of communication between him and me should be open."[252] One of the obstructions which may block a channel of communication can be uncooperativeness. In fact, Craig offers the description of objectivization in part as a way to ensure that his synthesized concept of knowledge includes subjects such as mobsters who are uncooperative and will not say that *p*, but nonetheless know that *p*. The need for an objectively good informant does not imply the need for a cooperative informant. How then may we use a lack of cooperativeness to explain why we are not willing to say that the wet street knows that it rained?[253]

iv. Using epistemic authority to make the distinction

With the epistemic authority hypothesis, we exclude considerations like convenience from the beginning as too subjectively idiosyncratic to affect the synthesis. We should not attempt to respond to this or any other objection by reverting to such subjective considerations. In section 3.a I point out that the epistemic authority hypothesis demands the explication of three new conditions: a social status condition, an honor condition, and a normativity condition. Any successful synthesis must show how the concept can meet these three demands of a social system of epistemic evaluation. Convention contextualism has been explicated specifically to enable the synthesis of these features together with sensitivity to practical environments. I now suggest that these three conditions indicate the proper way to distinguish between informants and mere sources of information, and thereby respond to Kelp's belief objection. The epistemic position of mere sources of information makes them into useful tools for finding out whether *p*, but fails to provide them with authority as to whether *p*.

In the context of a different discussion, Arnon Keren elaborates on the concept of epistemic authority, particularly its normative aspect.[254] According to his interpretation of testimonial interactions, an authoritative informant's assertion is not

[252] Craig 1990, p. 85.
[253] For another criticism of basing the distinction of convenience, see Feldman 1997, p. 209.
[254] Keren 2007. He concentrates on demonstrating that an informant needs to have epistemic authority in order for an inquirer to know that *p* on the basis of his testimony that *p*.

added to a bunch of reasons that the inquirer already has. Rather the informant's assertion *preempts* the other reasons the inquirer might otherwise have had to believe that *p*.

> To have the kind of epistemic authority that allows others to trust one on p, that allows others to take one's word for it, is to have a special kind of normative power, the power to issue a special kind of reason for belief: by expressing her belief that p, a person who has authority on p does not merely give us a reason to likewise believe that p. More than that, she gives us a second-order, preemptive reason for disregarding other relevant evidence which we may have concerning p.[255]

Keren points out further that the informant's preemptive power is justified by his responsiveness to reasons. I trust that Fred is responsive to adequate reasons, and *thereby* came about his belief that *p*. As a result, I do not need any further reasons in order to simply believe that *p* on his authority. If Fred has epistemic authority, then I trust that his reasons are good enough to justify his assertion. I believe him and in so doing, *make his reasons my own*. A subject who is not properly responsive to reasons does not have epistemic authority, because the normative aspect of authority implies that an inquirer ought to take the authoritative informant's reasons as her own.

Keren further explains epistemic authority as a matter of which inferences the inquirer makes in order to believe the informant's assertion.

> It is possible, in other words, to trust a speaker on the basis of an inference. But if one trusts a speaker, then one does not accept her judgment on the basis of an inference from what she has said to the fact that what she has said is true.[256]

An inquirer may need background information as to why a given subject is trustworthy, without altering his status as an informant. However, if she would need any *further* background information as to why his assertion that *p* implies that *p*, then his assertion could not be taken as a preemptive reason to believe that *p* and he should not be considered an epistemic authority on the question whether *p*.

Using Keren's definition of epistemic authority, we may explain why tree-climbing Fred is a good informant, while barn-gazing Henry is not. Once the inquirer hears Henry's assertion she must put it together with the other information she has about Henry's condition, and therefrom infer that Henry's assertion is true. She would not be able to simply take Henry's reasons for believing that *p* as her own reasons for believing that *p*, for she knows that he is not appropriately responsive to

[255] Ibid., p. 373.
[256] Ibid., p. 374.

reasons. Therefore, she may only take his assertion as one more piece of evidence, together with the rest of her information about his state.

On the other hand, once the inquirer has inferred that Fred is friendly, she trusts that he is appropriately responsive to reasons. Therefore, she may simply trust his assertion, assuming that all of his reasons for making the assertion are good enough to be her reasons as well. Even if Fred is entirely unaware of the yellow marauders, he is still in a good enough epistemic position to fully justify his assertion about the tiger, which is all the inquirer needs in order to trust him on that fact.

I do not pretend to have adequately supported Keren's proposal, nor am I even convinced that his account could overcome all objections. The work of finding the precise feature that really motivates this distinction requires more work and subtlety. However, I hope to have demonstrated the danger of falling into overly subjective considerations such as convenience when responding to fourth step discrepancies, and the ability of the epistemic authority hypothesis to do a better job.

Furthermore, we may go beyond Keren's work by using the *honor* condition to explain why Henry and other mere sources of information do not count as filling the relevant need. Ram Neta does just that, modifying Craig's hypothesis in response to the 'source of information' objection so that it describes the need for *creditable* informants.[257] Henry and his ilk do not deserve credit for their true assertions. They are being *used* by their inquirers, despite their major cognitive flaws. As Fricker recognizes, this is a degrading position to be in.[258] Only subjects who are appropriately responsive to reasons deserve the honor involved in a knowledge attribution. Thus the subjects who can be counted on to assert the truth *despite* their twisted responses to the relevant reasons do not qualify as knowers.

d. The scope objection

i. KAP–CC neglects theoretical knowledge

Throughout this thesis I have stubbornly ignored the question of skepticism. Given the motivation to explicate the concept of knowledge based on its actual social

[257] Neta 2006, p. 267.

[258] Fricker 2007, pp. 132 – 137. Fricker calls this relationship one of epistemic objectification, comparing it to sexual objectification. She concentrates on those cases in which it the objectification is wrongfully committed, whereas I concentrate here on how the 'mere source of information' label can be used for people like Henry, who deserve their objectification.

function, this omission may be acceptable. However, a closely related question cannot be entirely avoided: How can KAP–CC account for the use of the concept of knowledge when practical concerns are irrelevant?

Using KAP–CC requires reference to a hypothetical situation: in a nearby possible world, would the informant meet the demands of an inquirer's practical environment? However, there are many situations in which conventions may provide no clue as to which hypothetical practical environment we ought to take into consideration. A historian, for example, may happen to know the average price of cinnamon in medieval Portugal. What conventional standards govern that attribution of knowledge? In other words, which hypothetical practical environment do we intuitively judge the attribution against?

The scope of my synthesis seems to be too narrow. The concept of knowledge is surely used to cover lots of *useless* beliefs. Thus, it seems that very many uses of the concept of knowledge do not in fact function to play the role described by KAP–CC. Of course, there would still be the chance that KAP–CC describes one pattern of use among others, along Kusch's Wittgensteinian model.[259] However, the more ambitious hypothesis on which this synthesis has been based, that nearly *all* knowledge attributions serve the hypothesized function, would be untenable.

ii. Three possible responses

There are three ways in which theoretical uses of the concept of knowledge might indeed turn out to serve the hypothesized function. First, we could simply broaden the focus from conventional *practical* environments to *all* conventional environments. Conventional epistemic standards do not arise exclusively in response to practical considerations. There is also no reason to assume that we need to decide whom to trust only while practically reasoning; we are equally eager to keep our *theoretical* reasoning rational as well, and that involves just as much testimony and trust. In fact, Willaschek's epistemic practice contextualism does not mention practical reasoning at all.[260] He offers many examples of the kinds of standards that govern contexts such as natural science and the courtroom, which need not necessarily be connected to a hypothetical practical environment. Kappel also utilizes

[259] Kusch 2011. See section 2.b.iv.
[260] Willaschek 2007. See section 5.e.ii.

this more general form of contextualism in his practical explication, which results in the following condition:

> K(S_1, S_1-S_n, p) iff p, and S_1 is in a sufficiently good epistemic position such that S_1-S_n, given right circumstances of transmission, ought to take the truth of p for granted in their practical and theoretical deliberation.[261]

Kappel is neutral as to whether the relevant context of deliberation is theoretical or practical. Imitating Willaschek and Kappel, we may thus imagine the sorts of conventions governing whether or not one may rationally use *p* as a premise in one's *theoretical* environment, and thereby broaden the scope of KAP–CC to encompass more knowledge attributions.

This response involves very little further argument. Indeed, the only reason that I do not take such a line from the beginning, is that Hawthorne's practical environment is the most fully developed model with which we can distinguish the types of contexts to which knowledge attributions are sensitive. Elaborating a model of the specific factors that affect theoretical environments would have taken a great deal more constructive work, and it was more appropriate to the scope of this thesis to use the tools already available.

Despite the sufficiency of this moderate adjustment to quiet the objection, I would also like to suggest two other possible responses to the objection, both of which would require a great deal more argumentation, but which I leave here in the interest of peaking curiosity.

One more radical approach would be to broaden the notion of practical environments and of practical reasoning altogether. *Is* any knowledge entirely useless? The standards to which a subject holds her theoretical knowledge affect her epistemic character and self-image. It may well have far reaching and subtle consequences throughout her life. One might therefore render theoretic reasoning as a form of practical reasoning, carried out towards the intention of, for example, being a responsible person. Thus, one may see a historian's integrity as indeed a *practical* choice. She may count as knowing the price of cinnamon in medieval Portugal if someone could believe it on her authority and accordingly use it as a premise in the conventional practical environment of acting as a responsible person would act.

[261] Kappel 2010, p. 79.

The third available tactic would be to consider the alternative hypothesis that people need to explain the cognitive behavior of others. The concept of knowledge arguably functions to meet that need. One could then connect the need for labeling epistemic authority to the need to explain the cognitive behavior of others by means of a Williams-type imaginary genealogy. The genealogy would relate this *non-instrumental* explanation to its *instrumental* benefits for the individuals who use the social system of epistemic evaluation. One perk of such an alternative explication would be to more easily account for how knowledge functions when we attribute it to non-human animals, as ethologists often do.[262]

e. Suggestions for further research

i. Empirical data and normative stipulation

I pointed out that disagreements can arise as to which standards convention supports or ought to support. The question arises: What should count as a valid knowledge attribution, and what should count as a mistake? Such disagreement takes place within any society's actual epistemic practice and furthermore, were we to go on working with convention contextualism, it could just as well take place as we attempt to describe society's epistemic practice. I focus here on the tools that would be available to the latter group, the epistemologists working to clarify the conventional standards that govern the use of the concept of knowledge. I suggest that there are two avenues of clarification available. However, we should not assume that we will ever eliminate vagueness entirely. A good deal of epistemic practice does tolerate imprecision. There are no precise standards for petty bets among friends, and there arguably need not be.

Clarification can occur in two ways: empirical data and normative stipulation. Various attempts are being made to integrate epistemology with empirical studies. In

[262] As Hilary Kornblith convincingly argues in Kornblith 2002. Since knowledge can be applied to animals, Kornblith concludes in Kornblith 2011 that Craig's method is misled; the concept of knowledge is a reflection of a natural kind, so its extension cannot change relative to human concerns. His argument relies on the fact that biologists can definitively determine the environment relative to which an animal's cognitive process must be reliable to count as producing knowledge. Thereby, he overlooks the fact that *human* environments are not biologically determined, but rather irreducibly social. Thus, when we attribute knowledge to humans, social function does make a difference in the kind of reliability demanded.

relation to this work, doubt has been thrown on armchair analytic attempts to define knowledge. As Michael Bishop and J.D. Trout write,

> In fact, we find the temptation to construct a theory to separate out the 'justified' from the 'unjustified' oddly Scholastic – designed for a kind of prim conceptual tidiness rather than for useful guidance.[263]

The flexibility of KAP–CC provides a basic definition of knowledge which could be used and developed by empirical research. Furthermore, the synthesis could contribute to such research by suggesting fruitful directions. We must find empirical data to discern the allotment of epistemic authority in various fields of a given society and the ways these fields compare and interact. For example, a study of contrasting legal traditions can show various ways in which epistemic authority can be awarded to expert testimony.[264] A psychological study has revealed a difference between the indicator property [being a statistics professor] and the indicator property [being a psychology professor]: the former was generally taken to correlate more reliably with being right about the professor's own subject matter, whereas the latter was generally taken to correlate more reliably with being right about general life issues.[265]

Epistemological work on clarifying the concept of knowledge has relied largely on empirical *linguistic* data. When this data is seen through the model of KAP–CC, we find one important guideline: the examples of the use of the word 'knowledge' must not be underdetermined. The generated intuitions cannot be adequately interpreted if it is unclear which conventions affected the listener's understanding of the example. In addition to this guideline, the synthesis suggests several new questions that linguistic experiments should consider. Pairs of contrasting examples should be set up, so that we may better understand the difference in the standards implied by different conventional contexts. While the literature on contextualism has already begun this work, the choice of examples has been limited to a few stock favorites in which the subject does know in the first, but does not know in the second.[266] It would also be illuminating to test for subtler contrasts, with a wider range of conventional contexts. Also, in order to allow for recommendations and

[263] Bishop & Trout 2005, p. 116.
[264] Goldman 1999, pp. 304 – 311.
[265] Kruglanski, et al. 2005, p. 385.
[266] Blaauw 2005, p. 128.

social statuses, the public nature of epistemic authority requires the conventional context to be transmittable through second and third-hand reports. Examples should be constructed to study how these reports affect listeners. Again, while some work for and against contextualism has offered examples of such reports, they have been radically underdetermined. Finally, the neutrality of KAP–CC as to whether the conventional shift is semantic or pragmatic should encourage attention to the non-semantic features of case studies, such as the speaker's tone of voice.

Another way to deal with conventional standards is to normatively stipulate them. Indeed, when there is impassioned disagreement, this may well be the only way to achieve consensus. What should the standards of epistemic authority in education be? Should those standards change in a lesson on religion? How should the standards governing science be affected by those governing education, art, and religion? All these issues already inspire a great deal of work towards their resolution. From KAP–CC we may draw the lesson that the final answers need not be reductive. The goal is not to figure out the one appropriate set of standards of epistemic authority, for each field may and often must operate its own epistemic standards.

ii. Factual and conjectural genealogies

How does my use of the functional synthesis fit into the range of interpretations I describe in sections 2.b and c? As promised, I did not make use of any genealogy, but rather confined myself to the tools of Kappel's practical explication. However, by endorsing convention contextualism I also emphasize the incompleteness of the synthesized concept, KAP–CC, that has resulted. Kappel too ends up with a concept indexed to the needs of the user, but he does not pursue this sensitivity any farther. The omission is crippling, as such a conception of knowledge would fluctuate relative to individual rather than public contexts, and thereby lose authority and intelligibility.

The function of the concept of knowledge may only be fully comprehended when we understand the conventions actually at work in any given society. I have suggested that empirical sciences, particularly psychology, sociology and linguistics, may reveal the details of conventional contexts. Another way to broaden our empirical understanding of conventional contexts is to turn to history, i.e. to

construct a *factual* genealogy. A factual genealogy would illuminate the synthesis by showing how the use of the concept of knowledge has been modified by different conventions at different times and places. In addition, nothing bars the factual genealogy from including *conjectural* claims about the earliest uses of the concept of knowledge.

Furthermore, it is not necessarily an accident that I, and Craig, begin the synthesis with the most individualistic conditions and end with the conditions most involving society. It might be too difficult to begin by understanding the social aspects of the concept's function. In any case, it would be possible to rewrite the argumentation for convention contextualism as an imaginary genealogy. However, I have preferred to do without that expository device in order to maintain a maximum of clarity and simplicity.

7. Conclusion

The method of conceptual synthesis begins with the plausible function of a concept and, on that basis, explicates its conditions. While Craig pioneers the use of this method in epistemology, his own results remain stubbornly suggestive. His synthesis includes a genealogy, meant to illuminate the marks of the concept's function, but it is unclear just what those marks could be or what kind of genealogy should illuminate them. This fogginess has inspired some confusion, as several contradictory approaches all claim to employ Craig's genealogical method. I outline this array of receptions, but do not support any one kind of genealogy, preferring to skirt the controversy altogether and apply the method without using a genealogy at all.

I thus begin by claiming that the concept of knowledge *actually* fills the need I hypothesize. That need, like Craig's, synthesizes a goal and a strategy. The goal remains the same: pick out the subjects with the true beliefs. The strategy however is modified to accommodate the social and diachronic conditions in which any epistemic practice must be embedded. Thus, I do not hypothesize that inquirers need to detect subjects that could meet their own personal momentary epistemic needs. Rather, I hypothesize that inquirers need to recognize subjects who could meet the enduring epistemic needs of their entire community. I characterize this social diachronic need as the need for epistemic authority. We may explicate three further conditions from this new hypothesis. The concept which meets the need for epistemic authority would have to confer a recognized social status, honor, and a normative obligation to others, for those with epistemic authority *ought* to be believed.

Despite these three new conditions, I refuse to remove the individual momentary needs from the original hypothesis as well. Thus, I argue that the concept of knowledge functions to serve individual inquirers as they evaluate whether a belief would be trustworthy enough to use in their own practical reasoning. Any condition that function demands must accordingly be included in my synthesis. I find that this function demands sensitivity to the attributor's practical environment. Despite disregarding skepticism and focusing on practical reasoning, a realm typically conducive to subject-sensitive invariantism, my synthesis therefore includes a

contextualist condition. No other condition would render the synthesis flexible enough to describe the function of the concept of knowledge in scenarios of testimony evaluation, in which an inquirer expresses whether a potential informant would be reliable enough for her practical environment by attributing or withholding knowledge to him.

The sensitivity condition clashes with the diachronic needs of epistemic communities with which my synthesis begins. Specifically, it is difficult to see how an attributor-sensitive concept could maintain the intersubjective authority and intelligibility required by the hypothesis. If the explicated conditions cannot be synthesized into one coherent concept, then the hypothesis must be false. I investigate several conceptions of knowledge to see if any of them offer the synthesis my hypothesis would require. After several failed attempts, I show that convention contextualism offers a plausible solution. If practical environments are publically fixed in common conventions, sensitivity need not threaten intersubjective authority and intelligibility. Each context has its own absolute standard of authority and the pertinent context is nearly always detectable, even in passed down reports or overheard tidbits, thanks to a plethora of conventional cues. A brief survey of our intuitions regarding real knowledge claims shows that conventions do in fact play an irreducible role in how we understand their implications. Thus, with the help of convention contextualism, I present the final synthesis of the concept of knowledge. The synthesis includes a true belief condition as well as an epistemic authority condition, the standards of which are relative to the conventional context in which the concept is used.

The final chapter compares the synthesis with some widely-shared intuitions generated by the concept of knowledge, to ensure that there is no unacceptable discrepancy. I consider four objections which claim to pose trouble for the synthesis and argue against each of them. Despite some reasons to suspect the contrary, the synthesis can generate the appropriate intuitions in cases of disagreement and behavior evaluation. Furthermore, basing the synthesis on epistemic authority suggests a better means of distinguishing between knowers and mere truth-asserters than that provided by Craig's good informant based synthesis. Finally, I point out that the restriction in scope affected by my attention to practical environments need pose no grave problem either. While some adjustment is required to describe the

function of the concept of knowledge with respect to useless information, there are several plausible ways to do so.

In conclusion I point out that the proposed synthesis is well-suited for use as a working definition of knowledge in empirical and normative pursuits. Its open-ended sensitivity and social orientation make the synthesis a flexible and suggestive resource.

Bibliography

Adler, J. (2010 Edition, Winter). *Epistemological Problems of Testimony.* (E. N. Zalta, Ed.) Retrieved from The Stanford Encyclopedia of Philosophy: <http://plato.stanford.edu/archives/win2010/entries/testimony-episprob/>

Bach, K. (2004). *Knowledge in and out of context.* Retrieved Feb. 19, 2011, from Kent Bach's Home Page: http://userwww.sfsu.edu/~kbach/Knowledge&Context.pdf

Bach, K. (2005). The Emperor's New 'Knows'. In G. Preyer, & G. Peter (Eds.), *Contextualism in Philosophy: Knowledge, Meaning and Truth* (pp. 51 - 89). Oxford: Oxford University Press.

Beyer, C. (2007). Contextualism and the background of (philosophical) justification. *Grazer Philosophische Studien, 74,* 291–305.

Bishop, M., & Trout, J. (2004). *Epistemology and the Psychology of Human Judgment.* USA: Oxford University Press.

Blaauw, M. (2005). Challenging Contextualism. *Grazer Philosophische Studien, 69,* 127 - 146.

Brown, J. (2005, April). Adapt or die: The death of invariantism? *The Philosophical Quarterly, 55, No. 219,* 263 - 285.

Brown, J. (2005). Comparing contextualism and invariantism on the correctness of contextualist intuitions. *Grazer Philosophische Studien, 69,* 71 - 99.

Buckwalter, W., & Stich, S. (forthcoming). Gender and Philosophical Intuition. (September 26, 2010). Available at SSRN: http://ssrn.com/abstract=1683066

Cohen, S. (1988). How to be a fallibilist. *Philosophical Perspectives, 2, Epistemology*, 91-123.

Cohen, S. (2004). Knowledge, Assertion and Practical Reasoning. *Philosophical Issues, 14, Epistemology*, 482 - 91.

Craig, E. (1990). *Knowledge and the State of Nature.* Oxford: Oxford University Press.

Craig, E. (2007). Genealogies and the State of Nature. In A. Thomas (Ed.), *Bernard Wiliams.*

Crisp, T. (2005). Hawthorne on knowledge and practical reasoning. *Analysis, 65.2*, 138 - 140.

Davis, W. A. (2007). Knowledge claims and context: loose use. *Philos Stud*, 395 - 438.

DePaul, M. R. (2009). Ugly Analyses and Value. In A. Haddock, A. Millar, & D. Pritchard (Eds.), *Epistemic Value* (pp. 112 - 139(28)). Oxford: Oxford Scholarship Online Monographs.

DeRose, K. (2004, March). The Problem with Subject-Sensitive Invariantism. *Philosophy and Phenomenological Research, LXVIII, No. 2*, 346 - 350.

DeRose, K. (2009). *The Case for Contextualism.* Oxford: Oxford University Press.

Douven, I. (2008). Knowledge and Practical Reasoning. *Dialectica, 62*(1), 101 - 118.

Dretske, F. (1970, Dec. 24). Epistemic Operators. *The Journal of Philosophy, 67*(No. 24), 1007 - 1023.

Dretske, F. (1981). The pragmatic dimension of knowledge. *Philosophical Studies, 40,* 363 - 378.

Fantl, J., & McGrath, M. (2002, Jan.). Evidence, Pragmatics, and Justification. *The Philosophical Review, 111, No. 1,* 67 - 94.

Feldman, R. (1997, Mar.). Review Essay. *Philosophy and Phenomenological Research, 57*(No. 1), 205 - 221.

Fricker, M. (1998). Rational Authority and Social Power: Towards a Truly Social Epistemology. *Proceedings of the Aristotelian Society, New Series, 98,* 159-177.

Fricker, M. (2007). *Epistemic Injustice.* Oxford: Oxford University Press.

Goldman, A. (1999). *Knowledge in a Social World.* New York: Oxford University Press.

Greco, J. (2008, July). What's Wrong with Contextualism? *The Philosophy Quarterly, 58*(232), 416 - 436.

Greco, J. (2009). Knowledge and success from ability. *Philos Stud, 142,* 17 - 26.

Grundmann, T. (2008). *Analytische Einführung in die Erkenntnistheorie.* Berlin: Walter de Gruyter GmbH & Co. KG.

Hawthorne, J. (2004). *Knowledge and Lotteries.* Oxford: Oxford University Press.

Hawthorne, J., & Stanley, J. (2008). Knowledge and Action. *The Journal of Philosophy, 105*(10), 571-90.

Kappel, K. (2010). On Saying that Someone Knows: Themes from Craig. In A. Haddock, A. Millar, & D. Pritchard (Eds.), *Social Epistemology* (pp. 69-89).

Oxford: Oxford University Press.

Kelp, C. (2011). What's the point of "Knowledge" Anyway? *Episteme, 8*, 53-66.

Keren, A. (2007). Epistemic Authority, Testimony and the Transmission of Knowledge. *Episteme: A Journal of Social Epistemology, 4*(3), 368 - 381.

Kornblith, H. (2002). *Knowledge and its Place in Nature.* Oxford: Oxford University Press.

Kornblith, H. (2011). Why should we care about the concept of knowledge? *Episteme, 8*(No. 1), 38 - 52.

Kruglanski, A. W., Raviv, A., Bar-Tal, D., Raviv, A., Sharvit, K., Ellis, S., et al. (2005). Say's Who?: Epistemic Authority Effects in Social Judgment. *Advances in Experimental Social Psychology, 37*, 345 - 392.

Kusch, M. (2009). "Testimony and the Value of Knowledge". In A. Haddock, A. Millar, & D. Pritchard (Eds.), *Epistemic Value* (pp. 60-93). Oxford: Oxford University Press.

Kusch, M. (2011, Feb.). Knowledge and Certainties in the Epistemic State of Nature. (D. Pritchard, Ed.) *Episteme, 8*(No. 1), 6 - 23.

Lackey, J. (2007). Why we don't deserve credit for everything we know. *Synthese, 158*, 345 - 361.

Lackey, J. (2008). *Learning from Words.* Oxford: Oxford University Press.

Lane, M. (1999). States of Nature, Epistemic and Political. *Proceedings of the Aristotelian Society, 99*, 211 - 224.

Levin, J. (2008, March). Assertion, Practical Reason, and Pragmatic Theories of Knowledge. *Philosophy and Phenomenological Research, LXXVI*(2), 359 - 384.

Lewis, D. (1996, December). Elusive Knowledge. *Australasian Journal of Philosophy, 74, No. 4*, 549 - 567.

Menzies, P. (2009 Edition, Fall). *"Counterfactual Theories of Causation"*. (E. N. Zalta, Ed.) Retrieved March 18, 2011, from The Stanford Encyclopedia of Philosophy: <http://plato.stanford.edu/archives/fall2009/entries/causation-counterfactual/>

Neta, R. (2006). Epistemology factualized: New contractarian foundations for epistemology. *Synthese, 150*, 247 - 280.

Neta, R. (2006). Epistemology factualized: New contractarian foundations for epistemology. *Synthese, 150*, 247 - 280.

Owens, D. (2000). *Reason without Freedom: The Problem of Epistemic Normativity.* London: Routledge.

Richard, M. (2004). Contextualism and relativism. *Philosophical Studies, 119*, 215 - 242.

Rysiew, P. (2001). The Context-Sensitivity of Knowledge Attributions. *Noûs, 35:4*, 477 - 514.

Rysiew, P. (2007). Speaking of Knowing. *Noûs, 41:4*, 627 - 662.

Schaffer, J. (2005). What Shifts? Thresholds, Standards, or Alternatives? In G. Preyer, & G. Peter (Eds.), *Contextualism in Philosophy* (pp. 115 - 130). Oxford: Oxford University Press.

Sosa, E. (2005). A Defense of the Use of Intuitions in Philosophy. In M. Bishop, & D. Murphy (Eds.), *Stich and His Critics.* Blackwell Publishers.

Stanley, J. (2004). On the Linguistic Basis for Contextualism. *Philosophical Studies, 119 (1 - 2),* 119 - 146.

Steup, M. (2008 Edition, Fall). *The Analysis of Knowledge.* (E. N. Zalta, Ed.) Retrieved from The Stanford Encyclopedia of Philosophy: <http://plato.stanford.edu/archives/fall2008/entries/knowledge-analysis/>

Stich, S. (1990). *The Fragmentation of Reason.* Cambridge, Massachusetts: The MIT Press.

Unger, P. (1984). *Philosophical Relativity.* Oxford: Basil Blackwell Publisher.

Weinberg, J., Nichols, S., & Stich, S. (2001). Normativity and Epistemic Intuitions. *Philosophical Topics*(29), pp. 429-460.

Weiner, M. (2009). Practical Reasoning and the Concept of Knowledge. In A. Haddock, A. Millar, & D. Pritchard (Eds.), *Epistemic Value* (pp. 163 - 182). Oxford: Oxford University Press.

Willaschek, M. (2007). Contextualism about knowledge and justification by default. *Grazer Philosophische Studien, 74,* 251 - 272.

Williams, B. (2002). *Truth and Truthfulness.* Princeton, New Jersey: Princeton University Press.

Williams, M. (1996). *Unnatural Doubts.* Princeton: Princeton University Press.

Williams, M. (2001). *Problems of Knowledge.* Oxford: Oxford University Press.

Williamson, T. (2005). Knowledge, Context, and the Agent's Point of View. In G. Preyer, & G. Peter (Eds.), *Contextualism in Philosophy* (pp. 91 - 114). Oxford: Oxford University Press.

Williamson, T. (2007). *The Philosophy of Philosophy.* Malden: Blackwell Publishing Ltd.

Williamson, T. (forthcoming). Contextualism, Subject-Sensitive Invariantism, and Knowledge of Knowledge. *The Philosophical Quarterly*.

Wittgenstein, L. (2006). *Philosophische Untersuchungen.* Frankfurt am Main: Suhrkamp Verlag.

Abstract

Words and concepts are useful tools; they play a function vital to human survival. Conceptual analysis tends to ignore this fact, examining concepts in utter isolation from actual social life. As a result, conceptual analysts may unwittingly offer a definition that, were it correct, would leave the concept unfit to do the work it has to do. The method of conceptual synthesis, as developed by Edward Craig, corrects this mistake by explicitly constructing definitions so that they reflect the concepts' function. In the following thesis, I present Craig's conceptual synthesis of knowledge. I then elaborate on and amend his work with my own application of the method.

Craig begins with the hypothesis that having a concept of knowledge meets the need to evaluate whether an informant is trustworthy. This hypothesis fails because it does not fully account for the demands that social and diachronic circumstances place on the concept. I put forth an alternative hypothesis: having a concept of knowledge meets the need to label epistemic authority. Labeling epistemic authority involves recognizing a fixed social status and bestowing honor; therefore my hypothesis does not suffer from a disregard for social circumstances. Labeling epistemic authority is also important to subjects who are engaged in practical reasoning; they must evaluate which beliefs are trustworthy enough to use as reasons for action. To meet this practical need, knowledge attributions must mean different things to different people; the concept of knowledge must be relative to the practical context of the speaker. This contextual sensitivity may seem to interfere with the social function of knowledge. However the intersubjective authority and intelligibility of the concept are unaffected because the relevant contexts are constrained by public conventions.

People use the concept of knowledge in their practical reasoning as a label of epistemic authority and they interpret it as relative to conventional contexts. By constructing a definition that reflects this function, I offer a flexible, suggestive, and intuitively accurate conception of knowledge.

Abstract

Worte und Begriffe sind nützliche Werkzeuge; sie erfüllen eine dem menschlichen Überleben grundlegende Funktion. Begriffliche Analyse übersieht meistens diese Tatsache und untersucht Begriffe in totaler Isolation. Demzufolge bieten begriffliche Analytiker eine Definition, die, wäre sie richtig, den Begriff diese Funktion nicht erfüllen lassen. Die von Edward Craig entwickelte Methode der begrifflichen Synthese korrigiert diesen Fehler, indem sie Definitionen explizit erstellt, sodass sie sich an die Funktionen der Begriffen anpassen. In der vorliegenden Arbeit stelle ich Craigs begriffliche Synthese des Wissens dar und fahre fort, seine Arbeit mit meiner eigenen Anwendung der Methode zu ergänzen und verfeinern.

Craig beginnt mit der Hypothese, dass der Wissensbegriff das Bedürfnis erfüllt, andere anhand ihrer Vertrauenswürdigkeit einzuschätzen. Diese Hypothese scheitert, weil es die von sozialen und diachronischen Umständen geprägten Ansprüche an den Begriff nicht berücksichtigt. Ich stelle die Hypothese auf, dass der Wissensbegriff das Bedürfnis erfüllt, epistemische Autorität zu bezeichnen. Eine Bezeichnung der epistemischen Autorität schließt die Anerkennung einer gefestigten gesellschaftlichen Stellung und das Erweisen einer Ehre ein, weshalb meine Hypothese nicht an der Exklusion der sozialen Umstände leidet. Epistemische Autorität spielt auch für praktische Entscheidungen eine wichtige Rolle; Menschen müssen einschätzen, welche Überzeugungen vertrauenswürdig genug sind, um als Gründe für Handlungen verwendet zu werden. Um dieses praktische Bedürfnis zu erfüllen, müssen Wissenszuschreibungen verschiedene Bedeutungen in verschiedenen praktischen Kontexten haben. Die Kontextsensitivät scheint die soziale Funktion des Wissens zu beeinträchtigen, allerdings bleiben intersubjektive Autorität und Verständlichkeit des Begriffes unberührt, da die relevanten Kontexte durch öffentliche Konventionen fixiert sind.

Menschen verwenden den Wissensbegriff in ihren praktischen Entscheidungen als eine Bezeichnung der epistemischen Autorität und deuten ihn einem konventionellen Kontext entsprechend. Indem ich eine Definition erstelle, die diese Funktion richtig erkennt, biete ich eine anpassungsfähige, beziehungsreiche und intuitiv zutreffende Konzeption des Wissens.

Curriculum Vitae

Personalien

Martha Cassidy-Brinn
geb. 23.08.1985
California, USA

Ausbildung

2008 – 2011 Masterstudium der
Philosophie (Universität Wien)
2002 - 2006 Bachelorstudium
(Sarah Lawrence College)

Beruf

seit 2007 Lehrtätigkeit in öffentlichen
und privaten Bildungseinrichtungen