

Contextualism and the Neglected Question of Context

John Morrison
Williams College
May 13, 2001

Abstract:

A satisfactory contextualist theory of knowledge must provide an account of how knowledge varies across contexts. There are three contextualist proposals for developing such an account. This paper demonstrates that all of them are unacceptable. Contextualists have therefore failed to provide a satisfactory theory of knowledge.

Thesis written in partial fulfillment for honors in philosophy

Contents

Acknowledgements	iii
1. Introduction	1
2. A Selective History of Contextualism	5
3. DeRose's "Solving the Skeptical Problem"	10
4. Stalnaker and Pragmatic Presuppositions	21
5. Lewis' "Elusive Knowledge"	27
6. Concluding Remarks	33
7. References	35

Acknowledgements

During the fall of my junior year a Hampshire College philosopher under consideration for a position at Williams gave a talk on the relationship between epistemology and cognitive science. Afterwards I wrote an email to Professor White lamenting the fact that such an obviously gifted educator had devoted himself to, “the three headed hydra of analytic philosophy, the philosophy of mind, and cognitive science.” It is perhaps the greatest testament to Professor Cruz’s influence on my thinking that I now anticipate devoting myself to that very beast.

In my future writing and teaching, I hope to emulate the intellectual clarity and professorial dedication that Professor Cruz demonstrated as my advisor. One of my fondest memories is of sitting next to him on Christmas Eve as he patiently helped me craft a refined writing sample from a clumsy rough draft. That moment encapsulates the patience, diligence, and intelligence with which he guided my work. Whatever is redeeming in the following thesis is in no small measure due to his influence.

I am fortunate that Professor Cruz is only the latest in a long line of superior teachers. My parents’ influence on me has been immeasurable, inexpressible, and hopefully intractable. They blessed me with a happy childhood, and their selfless love emboldens me to this day. It was during dinner conversations about literature, history, computers, physics, politics, and religion that I learned to love books and big ideas. It is only now that I have come to appreciate all that they have done for me.

Following the eighth grade, I attended a debate camp run by a philosopher from Michigan State University. On the first day of camp, Professor Roper sent me to the library to read Rawls and Nozick. My first exposure to philosophy was so enjoyable that I decided to make a profession out of it. The next summer, in stroke of luck, I took “Reason and Argument” from Richard Gallimore, a graduate student at the University of Rochester. He widened my appreciation for philosophy, and introduced me to many of the questions that continue to engage me. His mentorship was instrumental in keeping me excited about ideas.

At Williams, the philosophy department provided me with a challenging and friendly intellectual environment in which to grow. Several classes and professors have been particularly influential. Professor Mladenovic’s introductory class gave me a sense of the rigor good philosophizing demands. Professor Dudley’s class on Kant’s *Critique of Pure Reason* was a rich introduction to systematic philosophy. I gain new insights into Kant each time I revisit the notes from his lectures. Professor White’s class on Heidegger’s *Being and Time* convinced me that phenomenology has something vital to contribute to the philosophy of mind. His openness to objections, and frequent use of examples, are qualities that I have tried to imitate in my writing and conversations.

Finally, I would like to thank everyone whose companionship returned me to this world whenever I lost myself amidst possible worlds. In particular, without Sean and Xavier I could not have endured writing this thesis, nor could I have withstood the agonies of the graduate school application process.

Most of all, I would like to thank Lia for tolerating my moods with love and patience. She is more than I deserve.

Introduction

Contextualism about knowledge is the view that what one knows depends upon what context one is in.¹ The contextualist motivates this claim by appealing to intuitions about language use. In order to clarify those intuitions, it is useful to begin with an analysis of the word “flat”.

A motorist from the Rocky Mountains might say that the roads in New York are flat. An engineer might say household mirrors are not flat while the mirrors on the Hubble Space Telescope are flat. A geometer might say that no object is flat because all molecular structures are uneven. We have strong intuitions that the motorist, engineer, and geometer have all made potentially appropriate claims. But one cannot say without contradiction that, “The roads in New York are flat but my mirror at home is not” or “The mirror on the Hubble Telescope is flat but no object is flat.” So can the motorist, engineer, and geometer all have made true claims?

Some might try to avoid the question by objecting that statements about flatness are too indeterminate to assign them truth-values. While this objection identifies a useful insight about “flat”, it errs by confusing vagueness for indeterminacy. In other words, it is correct that expressions about flatness are useless without any constraints on what surfaces count as “flat”. It is incorrect in assuming that it is impossible to constrain what surfaces count as “flat”. At least two interpretations of “flat” can address the problem of vagueness. The first is to stipulate that one standard for flatness applies in all cases. Typically, those inclined toward this view hold that the geometer’s definition of a flat surface is the only standard for flatness. While it sometimes might be convenient to assert that roads and mirrors are flat, those assertions are false since they cannot satisfy the geometer’s standard. The second interpretation allows facts external to the content of the assertion to fix the meaning of the term “flat”. In other words, the context fixes the meaning of “flat”. This is to treat “flat” as an indexical, not unlike “she”, “here” or “that”. Following Unger (1984) I call those who favor the first interpretation “invariantists”, and those who favor the second interpretation “contextualists”.

So, given that flat is not indeterminate, can the motorist, engineer, and geometer all have made true claims? One’s answer depends on whether one is an invariantist or a contextualist about the word “flatness”. An invariantist will claim that all three cannot have made true claims since there can only be one standard for flatness, and all three assertions depend upon different standards for flatness. A contextualist will claim that all three have made potentially true assertions because facts external to each utterance might have fixed the standard for flatness in such a way that each is true. In other words, each can be true within its respective context. The contextualist contends that just as the truth-value of “He is six feet tall” depends upon who “he” refers to, the truth-value of “This road is flat” depends upon what standard for flatness we are referring to.

¹ For an overview of the contextualist literature, see DeRose (1999) and Pryor (forthcoming).

A similar pattern develops for sentences containing the word “know”.² A child might claim to know the temperature by looking at a mercury thermometer. A meteorologist might claim to know that a mercury thermometer has given a correct reading only after he has checked to make sure the mercury column is unbroken. A skeptic might claim that he can never know the temperature because he cannot rule out the possibility that an omnipotent demon is fooling him. Assuming the child and meteorologist report the temperature accurately then we have intuitions that all three knowledge claims are appropriate.³ But as with the flatness example above, all three of these claims are incompatible. So can all three have made true claims with respect to knowledge? One’s answer depends on whether one is an invariantist or contextualist about knowledge. The invariantist denies that all three claims to know can be true because each claim depends upon a different standard for knowledge. The contextualist thinks it is possible for all of the claims to be true because she allows for variances in the standard for knowledge.

We are confronted with apparently contradictory intuitions about knowledge. On the one hand, it seems that knowledge is bimodal: someone either knows *P* or does not know *P*. Assuming one finds the skeptic’s arguments compelling, then one must conclude that no one can truthfully ascribe knowledge. On the other hand, outside of philosophical contexts, it is intuitively appropriate to ascribe knowledge to others and to ourselves. In a courtroom, on a construction site, or during a telemarketing call, we would consider it quite strange if someone cited the possibility of deception by an omnipotent demon as a reason for refusing to ascribe knowledge.

The contextualist argues that her interpretation of knowledge satisfies both intuitions. It satisfies the first intuition by allowing knowledge claims to have a determine truth-value relative to a specific context, and it satisfies the second intuition by maintaining that the standards for knowledge vary across contexts.

The invariantist rejects this. She thinks that the contextualist has misinterpreted the term “know”. For instance, some invariantists accuse the contextualist of confusing assertibility conditions with truth-conditions (Unger, 1975). Assertibility conditions differ from truth-conditions because sometimes it is useful to speak falsely. The contextualist responds that by default we should interpret our language in whatever way makes the greatest number of our sentences true (DeRose, 1999). Other invariantists claim that knowledge is not an indexical, and there is no evidence that truth-conditions vary for non-indexical terms (Schiffer 1995, Stanley 1999). One reason for thinking this is that if someone says to you “He knows that it is winter” and you ask the speaker who “he” refers to, the speaker should have little trouble providing the missing information. However, if you ask the speaker what standard for knowledge “knows” refers to, you are unlikely to elicit anything more than a blank stare. The contextualist counters by claiming that knowledge is a hidden-indexical (Hoffweber, 1999). A hidden-indexical is a phrase whose referent is not fully understood by the speaker. For instance, if a Roman soldier said the Latin equivalent of “it is winter” while invading the British isle during December, we would take him to have spoken a true sentence even if he was unaware that it is always the opposite season in the southern hemisphere.

² See Unger (1975), Dretske (1981), and Lewis (1996).

³ The contextualist takes as her starting point the fact that people find skeptical arguments compelling. The historic preoccupation with skepticism, and the contemporary dissatisfaction with proposed solutions (see Stroud 1984), makes the contextualist’s starting point reasonable.

It remains an open question whether contextualists or invariantists offer a superior interpretation of the word “knowledge”. The contemporary literature is almost exclusively concerned with developing this debate.

This paper puts aside the question of interpretation, and examines contextualist proposals for explaining *how* knowledge varies across contexts. A contextualist must be able to explain why the child has knowledge without checking for a broken mercury column or disproving the evil demon hypothesis, why the meteorologist must check for a broken mercury column but need not disprove the evil demon hypothesis, and why the skeptic can ignore neither the possibility of a broken mercury column nor the evil demon hypothesis. If the contextualist is to provide a satisfactory theory of knowledge then she must be able to explain what causes these shifts in the standard for knowledge. That is, she must provide an account of how contexts determine the relevant standard for knowledge.

It is therefore peculiar that attempts to define “context” in the current contextualist literature are virtually nonexistent. After all, if a utilitarian failed to provide a definition of utility then we would withhold judgment until one was provided. Similarly, if a foundationalist failed to specify what the first element of the structure of justification was, we would think that her theory was incomplete. The failure of contextualists to explicitly define “context”, or specify its relation to knowledge, makes one worry that contextualism is likewise an incomplete theory.

Contextualists do explicitly answer the question of *whose* context is relevant in a knowledge ascription. Contextualists make facts about both the subject and attributer’s context relevant. This distinguishes contextualists from those who think knowledge is relative to only the subject’s context. However, to say simply that knowledge is relative to facts about two contexts is to say little. One must specify which facts.

Contextualists sometimes list a few of the relevant facts about contexts in an attempt to demonstrate that contextualism can handle epistemology’s traditional problem cases, including Dretske’s (1970) zoo case, Harman’s (1973) newspaper case, and Ginet’s barn case (Goldman 1976). However, when they do so, they seldom provide general principles that are applicable to other cases as well. The relation between context and knowledge is left ambiguous. For instance, Stine (1978) thinks that evidence determines the probability of all alternatives, and the conversational context of the ascriber determines how likely an alternative must be in order to be relevant to ascription. Stine thinks that if a possibility is contextually unlikely enough, the putative knower does not need evidence to eliminate it (258). While this provides for a potential response to the skeptic she admits that it leaves, “rather indeterminate the way in which the context determines what is to be taken to be a relevant alternative” (249). Dretske (1981) only provides “some of the considerations” that affect what is relevant for knowledge ascription (373). The considerations he does provide are self-admittedly vague. For instance, he stipulates that one should determine whether or not a counterfactual possibility is relevant based upon, “the kind of possibilities that actually exist in the objective situation” (377). Cohen (1988) is more specific, identifying four contextual components: the subject’s evidence, what possibilities the subject is thinking about, counterfactual likelihood, and the attributer’s standards for knowledge at the moment of attribution. However, he acknowledges that he is unable to explain how these factors operate, or how they interrelate (see Cohen 1991, 32-34; Cohen 1988, 115-117). DeRose (1992) provides three factors that can contribute to making a possibility relevant,

including the mentioning of it, the consideration of it, and the importance that one's attribution is correct (914-915). He does not work out any of these criteria, or explain how they interrelate.

One might therefore wonder why contextualism has become one of the most "widely discussed" theories of knowledge (Pryor, forthcoming). The reason is that contextualists have focused their efforts on providing a response to skepticism. Contextualists argue that we have two competing intuitions about skeptical arguments. The first is that the skeptical argument is compelling. The second is that we have knowledge in mundane cases without having to eliminate skeptical possibilities. The contextualist argues that the very existence of these intuitions implies that there are at least two kinds of contexts: those in which skeptical hypotheses are relevant, and those in which they are not. Therefore, the contextualist believes that, if she confines herself to answering skeptical challenges, that she can leave the concept of context undetermined (see Cohen 1991, 33).

However, if the contextualist leaves herself free to define the context in whatever way is expedient for the case she is presently considering, then one wonders if she is doing anything more than reporting that fact that we do claim to have knowledge, and providing *ad hoc* theories in support of this observation. This in turn casts doubt upon her claimed ability to solve the skeptical problem. As long as contextualists fail to specify how knowledge relates to context in a generalized way then the entire theory remains "unacceptably occult" (Sosa 1986, 585).

In the current literature one can find three suggestions for dealing with this problem. None are developed at any length, and two must be extracted from the ways in which contextualist authors sometimes handle problem cases. This paper identifies each strategy, develops it, and evaluates it. In the end, it rejects them all. Therefore, no current contextualist account of knowledge is satisfactory.

The following section provides a historical introduction to contextualism. The third section examines DeRose's definition of contexts in terms of intuitions about possible worlds in "Solving the Skeptical Problem". The fourth section looks at a discarded suggestion from Goldman (1966) and Stalnaker (1972) to make knowledge relative to the presuppositions of a speaker. The fifth section evaluates Lewis' dependence upon the concept of conversational scorekeeping in "Elusive Knowledge". Concluding remarks are made in the sixth section.

A Selective History of Contextualism

Philosophical lineages are notoriously difficult to trace. Contextualism is no exception. One could select Wittgenstein (1969), Unger (1975), Dretske (1970), or Goldman (1976) as the founding contextualist figure and then trace the evolution of contextualism in a number of different directions. I have chosen one pedigree over others for a practical reason. I wish to explain why contextualists make the meaning of the word “know” depend on the context. Towards this end I trace the evolution of contextualism from Dretske (1970) to Stine (1976) to Cohen (1988).

While Dretske (1970) is not a contextualist paper, it established the framework within which contextualism emerged. He argues that knowledge is a semi-penetrating operator. An operator F is penetrating if the following syllogism necessarily holds:

- (1) $F(x)$
- (2) $F(x \text{ therefore } y)$
- (3) $F(y)$

Examples of penetrating operators include “it is true that”, “it is a fact that”, “it is necessary that”, and “it is possible that” (1007). An operator is semi-penetrating when the above syllogism holds in some cases, but not in all cases. Drawing upon his 1969 book *Seeing and Knowing*, Dretske claims that seeing is a semi-penetrating operator. For instance, “To see that the widow is limping is not to see that it is a widow who is limping” (1013). Like seeing, Dretske claims that “all epistemic operators are semi-penetrating operators” (1009). This means “we can say that S knows that X is Y without implying that S knows that *it is* X which is Y ” (1013).

Dretske provides two sets of examples to show that knowledge does not penetrate. The first set of examples is intended to be uncontroversial and offer support to the more controversial second set of examples. An example of the first type is as follows (see 1013). Dretske’s brother boards a bus in a city and takes a seat. An old lady then boards the bus and, finding all the seats occupied, begins to glower at Dretske’s brother in an attempt to make him relinquish his seat to her. After several minutes she gives up and walks to the back of the bus. Dretske argues that if knowledge were a penetrating operator the following would hold:

- (4) The little old lady knew that Dretske’s brother would not move
- (5) The little old lady knew that Dretske’s brother’s not moving entails that it was Dretske’s brother who would not move
- (6) The little old lady knew that it was my brother who would not move

Clearly (6) is false since the little old lady didn’t know who Dretske’s brother was. Dretske therefore claims that (5) is false. He writes, “We can say that she [the little old lady] realized that, or knew that, *my brother* would not move (minus, of course, this

pattern of emphasis), and we can say this because saying this does not entail that the little old lady knew that, or realized that, it was my brother who refused to move” (1013). Dretske uses this example as proof that knowledge is a semi-penetrating operator.

The second type of example is as follows. Imagine that you take your son to the zoo. You see what appear to be zebras. Your son asks if they are zebras, and you respond that you know that they are zebras. But if they are zebras then they necessarily cannot be mules cleverly painted to look like zebras. However, all of your evidence is visual and if they are cleverly painted mules that look like zebras then your evidence is of no help when deciding between the two hypotheses. Thus, the possibility that they are mules cleverly disguised to look like zebras is consistent with your evidence. This allows the skeptic to offer the following argument:

- (7) I know that those animals are zebras.
- (8) I know that those animals being zebras entails that those animals are not mules cleverly disguised to look like zebras
- (9) I know that those animals are not mules cleverly disguised to look like zebras.

(9) is false because all your evidence is consistent with the animals being cleverly disguised mules. Thus either (7) or (8) is false. The skeptic wants us to deny (7). Dretske wants to maintain our intuition that (7) is true so he must deny (8). In the crowded bus example Dretske takes himself to have proven that knowledge is a semi-penetrating operator. Dretske argues that if knowledge is a semi-penetrating operator, then we can reject (8). Dretske provides no criterion for determining when knowledge penetrates and when it does not, but expresses certainty that knowledge will never penetrate for skeptical hypotheses.

Stine (1971) rejects Dretske’s argument on the grounds that (4) is a false inference from Dretske’s first example. Stine argues that we should replace (4) with:

- (4’) My brother is such that it was known by the little old lady that he would not move.

When we replace (4) with (4’) then the argument becomes invalid. Thus Dretske has not proven that knowledge is a semi-penetrating operator. Since the case of the old lady case was the justification for his treatment of the painted mule case, Dretske has not given us grounds for rejecting (8) and retaining (7). Dretske has therefore failed to provide us with a response to skeptical challenges that maintains the truthfulness of mundane knowledge ascriptions.

Although Stine (1971) rejects Dretske’s answer to the skeptic she sees in Dretske (1970) a promising new tactic for responding to skepticism. Since both Stine and Dretske ultimately wish to develop a response to the skeptic, it is helpful to set up the skeptic’s argument in a general form.¹

Let H be a skeptical hypothesis and O be a proposition that one would ordinarily claim to know. So, for instance, O might be “I am sitting in a chair by a fire” and H might be “I am being fooled by an omnipotent demon”. Alternatively O might be “I am

¹ This formulation comes from DeRose (1995). DeRose’s skeptical argument is constructive, which is more useful for our purposes than Dretske’s (1970) argument by contradiction.

seeing zebras” and H might be “I am seeing painted mules”. The skeptic argues that the following syllogism holds:

- (10) I don't know that not-H;
- (11) If I don't know that not-H, then I don't know O;
- (12) Therefore, I don't know that O.

If H is well chosen then there can be no evidence that not-H. For instance, any evidence that I am not being fooled by an omnipotent demon might have been created by an omnipotent demon. Likewise, any evidence that I am not dreaming might merely be a product of my dream. If the skeptic chooses H properly, and one accepts both of her premises and her logical inference, then the skeptic has demonstrated that we know nothing at all. The skeptic could also exploit more specific skeptical hypotheses. For instance, as we saw in the zebra case, H can be chosen so that it undermines a limited class of knowledge claims. Such hypotheses exploit the fallibility of evidence.

Stine wishes to deny (12). Moreover, she argues in her (1976) that we can resist the skeptic without denying that knowledge is a penetrating operator. She therefore accepts (11). Since she thinks the skeptical argument is valid, consistency demands that she deny (10). However, she wants to reject (10) without resorting to the dogmatic strategy of Moore (1962), who thought it sufficient to insist that he knew he had two hands regardless of the skeptic's tricks.

Stine's strategy for denying (10) is to claim that we know not-H when not-H is not a relevant alternative. For example, Stine claims that the possibility you are seeing a painted mule is not a relevant alternative, and therefore you do not need evidence to know that you are not seeing a painted mule. Stine's burden is to identify criteria for determining what alternatives are relevant at a given moment.

According to Stine there are two ways to define a relative alternative. First, one can identify an alternative as relevant “only if there is some reason to think it *could* be true”(252). Dretske's argument depends on this definition of a relevant alternative. He supposes that an alternative is relevant if there is no evidence to contradict it. One can determine the set of relevant alternatives by considering all mutually exclusive propositions, and eliminating only those that are logically incompatible with one's evidence. Since we cannot have evidence that directly contradicts the omnipotent demon hypothesis, the omnipotent demon hypothesis is always a relevant alternative. If one adopts this definition of a relevant alternative, then, according to Stine, the only way to retain one's anti-skeptical intuitions about knowledge is to follow Dretske and deny that knowledge is a penetrating operator.

One can also classify an alternative as relevant “only if there is some reason to think that it is true”(252). Stine adopts this second definition. For Stine, the phrase “some reason” is equivalent to the phrase “some evidence”. Therefore, for an alternative to be relevant there must be sufficient evidence to think it is true.

When does one's evidence give one sufficient reason to think that an alternative is true? The standards change depending on what context one is in (254). Stine admits that she is unable to be specific about how the standards for evidence vary with contexts (249, 252). However, Stine thinks that because we have no evidence that an omnipotent

demon is fooling us, in most contexts that possibility is not relevant for knowledge.² Likewise, in the zebra case, one usually does not have sufficient evidence to believe that one is really seeing a painted mule and therefore this possibility is irrelevant for knowledge.

If an alternative is not relevant then one knows that it is false. Stine writes, “[I]f the negation of a proposition is not a relevant alternative, then I know it – obviously, without needing to provide evidence ...”(1976, 258). Applying this to the zebra case, one can usually know that one is not seeing a painted mule without having any evidence to that effect.³ Likewise, if in a particular context the possibility that you are being fooled by an omnipotent demon is not relevant then you can know it is false without any evidence.

The closure principle of knowledge is the principle that knowledge is closed under logical entailment. For instance, if S knows P, and P implies Q, then according to the closure principle, S knows Q. In Dretske’s terminology, the closure principle is the principle that knowledge is a penetrating operator.

Stine’s paper intends to show that Dretske need not sacrifice the closure principle of knowledge in order to refute the skeptic. By claiming that we usually know that skeptical hypotheses are false without evidence, Stine responds to the skeptical argument by denying (10). This allows her to accept (11) and therefore preserves closure.

Cohen (1988) argues that Stine’s attempt to preserve the closure principle of knowledge is unsuccessful.⁴ Cohen points out that Stine is committed to denying that having sufficient evidence is a penetrating operator. Stine denies this principle by claiming that we do not need evidence to know not-h if not-h isn’t a relevant alternative (see 258).

It is helpful to think of this in terms of Dretske’s original syllogism. Let K be the operator “knows that” and E be the operator “has sufficient evidence to know that”. Let H be a skeptical hypothesis and O be a proposition one would normally claim to know.

(13) K(O)

(14) K(O therefore not-H)

(15) K(not-H)

(16) E(O)

(17) E(O therefore not-H)

(18) E(not-H)

Stine claims that we should accept (15) but reject (18).

Stine argues that she maintains closure by “giving up other less obvious principles” (251). Stine’s “less obvious” principle turns out to be the closure of evidence.

Stine demonstrates that giving up the closure principle for knowledge is undesirable, but does nothing to prove that giving up the closure principle for evidence is less undesirable. There are two good reasons for thinking that evidence and knowledge must be the same kind of operator. As Cohen (1988) points out, many people believe that

² Stine leaves it an open question whether or not there are contexts in which we have reason to believe the omnipotent demon hypothesis is true (257).

³ Unless, of course, they are in fact painted mules. But knowledge traditionally has a truth condition to handle exactly this sort of case.

⁴ Cohen also offers a few suggestions for developing a contextualist account of knowledge. However, these suggestions are worked out in greater detail in DeRose (1995), and therefore I will delay addressing them until the next chapter.

skeptical doubts are really doubts about evidence. Rephrasing the skeptical debate in terms of evidence offers such people no advantage.

We may add that most people think knowledge is at least a partial function of evidence. Stine asks such people to believe that knowledge and evidence aren't inextricably linked. Stine writes of Dretske, "Dretske is deluded by the fact that many knowledge claims require evidence on the part of the knower into thinking that all knowledge claims require evidence" (257). But requiring Dretske and others to give up their belief in the absolute relation between evidence and knowledge necessitates violating their intuitions about evidence and knowledge. Stine justifies this violation of intuitions by distinguishing between using a concept truthfully and using a concept correctly (see 255-256). Stine's claim is that it is often appropriate to assert that knowledge requires evidence, but that this does not change the fact that knowledge does not require evidence. This invokes the same distinction that the invariantist about knowledge relies on to explain our intuition that we have knowledge in mundane situations. It is therefore hard to imagine how she can make this distinction while remaining a contextualist about knowledge. If Stine is going to take some of our intuitions seriously, and not others, then she needs to provide arguments.

Contextualists wish to deny (12). This implies they must deny (10), (11), or the validity of the argument. Dretske's and Stine's failures convincingly demonstrate the problems with denying (10) or (11). Contextualists choose to deny the validity of the argument. One can safely assume that Dretske's and Stine's failure to convincingly deny (10) or (11) influenced this decision.

Notice that (12) follows from (10) and (11) only if the meaning of "know" is constant for (10), (11), and (12). The contextualist strategy is to provide an account of knowledge that indexes the meaning of a knowledge claim to a context, and then claim that while (10) and (11) are true, they are true relative to different contexts.

How is a knowledge claim relative to a context? The contextualists build upon the tradition of Dretske (1970) and Stine (1976) by claiming that the meaning of knowledge claims are relative to a set of relevant alternatives. They differ from their progenitors by arguing that the *context* determines what counts as a relevant alternative. For instance, the contextualist claims that in most contexts the omnipotent demon hypothesis is irrelevant. Therefore, even if we do not know that it is false, we can still know the time of day, where one's hometown is located, and who won the Super Bowl. The omnipotent demon hypothesis is relevant in some contexts, like epistemology classrooms. In those contexts the skeptical argument becomes valid and we therefore do not know anything. The contextualist makes what we know depend on what alternatives are relevant in the present context.

This preserves the closure principle of knowledge by making knowledge fully penetrating for a set of relevant alternatives. In other words, knowledge does not have to penetrate to irrelevant alternatives because, when we attribute knowledge, we mean that someone knows relative to a particular set of relevant alternatives. If someone claims that knowledge should penetrate to otherwise irrelevant alternatives then they mean something different by "know".

How does the context determine what is a relevant alternative? Contextualists seldom directly answer this question, preferring to focus on issues surrounding the contextualist interpretation of "know". The remainder of this paper criticizes the four

contextualist proposals for clarifying how the context determines what is a relevant alternative.

DeRose's "Solving the Skeptical Problem"

Many give credit to Plato for identifying the three traditional criteria for S to know P: that P is true, S believes P, and S's belief is justified. Regardless of its actual historical origin, philosophers for the last two millennia tended to take this definition for granted, assuming that if they could identify what is required for the justification criterion then they would have a full understanding of what is required for knowledge.

In 1963 Edmund Gettier published "Is Justified True Belief Knowledge?", challenging the tradition that Plato allegedly began.¹ Gettier identified cases where people formed beliefs through proper reflection on the evidence, and hence formed a justified belief, but nonetheless have true beliefs only by accident. For instance, suppose Smith and Jones are co-workers, and their boss announces one morning that he will promote one of them that afternoon. Suppose that Smith talks with the boss before lunch, and his boss assures him that Jones will receive the promotion. Smith consequently forms the justified belief that Jones will receive the promotion. Smith also justifiably believes that Jones has ten coins in his pocket; perhaps Smith saw Jones count the coins as he placed them in his pocket, and heard him comment that he would keep them there all day. Smith consequently forms the justified belief that the person who will receive the promotion will have ten coins in his pocket. Smith is unaware that he has ten coins in his own pocket, and that Smith is the one that is going to receive the promotion. Smith's belief that the person with ten coins in his pocket will receive the promotion is therefore true. Thus, Smith has a justified true belief. But does Smith know that the person with ten coins in his pocket will receive the promotion? Most epistemologists claim that he does not have knowledge because his justified true belief is only accidentally true; if Smith did not have ten coins in his pocket, then Smith would still believe the same thing.

In response to this kind of case, some epistemologists began to require that, in order for S to know P, P must cause S's belief that P (Goldman 1967). Smith's belief does satisfy this requirement because the fact that Smith is the one promoted, and the one with ten coins in his pocket, is not responsible for causing his belief that the man with ten coins in his pocket will receive the promotion.

In critically assessing causal theories of knowledge, including his own (1967), Goldman (1976) argues the fact that p might accidentally cause S's belief that p. Therefore, even if the fact that p causes S's belief that p it does not follow that S's belief is justified. For instance, suppose that Jones does receive the promotion, but there is a large hole in his pocket, and it is only by a stroke of luck that a coin does not fall to the floor.² Does Smith still know that the person with ten coins in his pocket will receive

¹ Russell (1912) contains a similar argument to Gettier (1960). However, the argument did not achieve widespread recognition until Gettier (1960).

² Goldman (1976) bases his analysis on the fake barn case. In order to streamline the discussion, I have modified Gettier's (1960) promotion case to make the same point as Goldman's fake barn case.

the promotion? Many would deny that he does because luck played too large a role in his being right.

Goldman concludes that belief *p* needs to exhibit sensitivity in counterfactual scenarios in order for belief *p* to be justified. In *Philosophical Explanations* Nozick formulates Goldman's insight as, "A belief that *p* is knowledge that *p* only if it somehow varies with the truth of *p*" (Nozick, 1981, 208). Nozick explains this variance in terms of possible worlds.³ According to Nozick, *S* knows *p* only if *S*'s belief that *p* matches the fact of the matter in close worlds. He calls this a sensitivity requirement. Applying this to our present example, Smith does not know that the person with ten coins in his pocket will receive the promotion because in close worlds where a coin falls out of Jones' pocket, Smith's belief will not reflect this change.

DeRose (1995) aims to merge Nozick's insistence on belief sensitivity with the contextualist's insistence that the standards for knowledge vary across contexts (33). He claims to accomplish this in a new criterion for knowledge that he calls The Rule of Sensitivity. Combined with a truth-condition, The Rule of Sensitivity is supposed to replace the traditional justified-true-belief account of knowledge. This section argues that DeRose's attempt to build a contextualist theory of knowledge on a Nozickean analysis of sensitivity is unsuccessful. Moreover, the analysis generalizes to show that contextualists cannot rely upon a sensitivity requirement to explain the contextual dependence of knowledge.

The Rule of Sensitivity is follows: "When it's asserted that *S* knows (or doesn't know) that *P*, then, if necessary, enlarge the sphere of epistemically relevant worlds so that it at least includes the closest worlds in which *P* is false" (37). In order to capture Nozick's tracking conditions, DeRose requires *S*'s belief to match the facts of the first not-*P* world, and all closer worlds.

All possible worlds are at different distances from the actual world. According to DeRose, worlds where skeptical hypotheses are true are very distant, and worlds similar to our own are very close. Thus, "distance", means something akin to "degree of similarity". By the Rule of Sensitivity, when any world becomes relevant then, necessarily, all closer worlds become relevant.⁴ In order to simplify the discussion, I will often refer to "possible worlds models". Each possible worlds model is a unique assignment of distances to all of the possible worlds. Thus, two possible worlds models differ if they place the same possible world at different distances. We will put aside for a moment the question of how DeRose intends to determine the distance to each world.

DeRose provides the following example:

³ Possible worlds present a way of understanding the truth-conditions of modal claims, such as "It is possible for Al Gore to win the next presidential election" and "it is possible for Al Gore to have won the last election" (see Samuel, Guttenplan 1994, 484). There is disagreement about the ontology of possible worlds, but neither Nozick (1984) nor DeRose (1995) require one to commit either way. For both of them, possible worlds are just convenient ways of talking about modal claims. For examples of two view on the ontology of possible worlds, see Lewis (1986) and Stalnaker (1984).

⁴ In an unfortunate, but common, overuse of the metaphor "world", DeRose thinks of possible worlds as occupying orbits around our own, not unlike the way the planets revolve around our sun at different distances. This is probably why he refers to the "*sphere* of epistemically relevant worlds" in the Rule of Sensitivity (37, emphasis mine). One will have much less trouble accepting talk about possible worlds if one instead thinks of possible worlds as sets, and distances as similarities between sets.

Consider my belief that I have hands. I believe this at the actual world, and it's true. What's more, in the other nearby worlds in which I have hands, I believe I do. There are also, at least in my own case, some alarmingly close worlds in which I don't have hands. These include worlds in which I lost my hands years ago while working on my uncle's garbage truck. In the closest of these not-P worlds, I'm now fully aware of the fact that I'm handless, and my belief as to whether I have hands matches the fact of the matter. My belief as to whether I have hands doesn't match the fact in various worlds in which I'm a BIV [brain in a vat], of course, but these are *very* distant (34).

Whether or not DeRose knows he has two hands hinges upon whether or not distant BIV worlds are relevant. In most contexts, only very close worlds are relevant, so DeRose knows he has two hands. In order for the skeptic to be successful, she must somehow make BIV worlds become relevant. As soon as BIV worlds become relevant, DeRose's belief that he has two hands will no longer vary with the facts in all relevant worlds because he would still believe he has hands if he was a brain in a vat. It therefore would not satisfy the Rule of Sensitivity. How could a skeptic change what worlds are relevant? One way is by mentioning skeptical possibilities. Consider the following exchange between Nate and his skeptic friend Sean.

- (1) Nate: "I know that I have two hands."
- (2) Sean: "Do you know that you're not a brain in a vat? If not, how can you claim to know you have two hands?"
- (3) Nate: "You are right, I don't know that I have two hands."

According to the contextualist, both of Nate's knowledge ascriptions are true. This would seem to be a contradiction, but the contextualist counters that the word "know" means two different things in (1) and (3).⁵ What has happened to cause this shift in meaning? Sean the skeptic has introduced the possibility of Nate being a brain in a vat, expanding the set of relevant worlds to include distant worlds. In this new enlarged set of relevant worlds, Nate's self-ascription is not sensitive since he might still believe he had two hands if he were a brain in a vat. Therefore, Nate no longer knows that he has two hands.

Likewise, in the painted mule case from the second section, you ordinarily have sufficient evidence to believe that you aren't seeing painted mules based upon your visual evidence and past experiences at zoos. If, on the other hand, someone were to press the issue and you began to seriously wonder if what you saw were zebras, it might not be true that you continue to know you are seeing zebras. This is because a possible world where zookeepers secretly paint mules to look exactly like zebras might have become relevant, and by hypothesis, your belief would fail to match the facts of such a world. This does not contradict your earlier claim to know that they are zebras because the meaning of both claims is relative to different sets of possible worlds.

This makes clear how developments in a conversation can result in a shift in the meaning of the word "know". It is also apparent that, for DeRose, the "context" of a knowledge claim is nothing more than the set of currently relevant worlds. Shifts in context are just shifts in the number of counterfactual worlds where one's belief must

⁵ See DeRose (1992, 924-928) for a more detailed defense of this claim.

match the facts. One can also think of a "context" as the distance that divides relevant worlds from irrelevant worlds.

Notice that no skeptic will be content with this analysis. Consider again DeRose's claim to know that he has two hands. The skeptic will demand that DeRose provide evidence for the claim that he lost his hands in a garbage truck accident in a close world while he is a brain in a vat only in a very distant world. But if DeRose concedes to the skeptic's demands then he is worse off than when he started. He would then have to prove that the chances of his being a brain in a vat are less than the chances of his having lost his hands while working on his uncle's garbage truck. Proving that he is not a BIV in close worlds requires him to prove that he is not a BIV in the actual world. As pointed out previously, skeptical doubts are at heart doubts about evidence. Skeptics try to demonstrate that our evidence, no matter how comprehensive, does not entail that skeptical hypotheses are false. DeRose concedes that, when the skeptic makes his argument, we must admit that he is correct. This entails conceding that we do not have evidence to rule out skeptical alternatives. Thus, DeRose's reliance on a possible worlds model potentially begs the skeptical question in dispute.

DeRose is aware of this problem:

Along the way, I've been assuming certain things that we believe but that the skeptic claims we can't know, thereby perhaps raising the concern that I'm begging the question against the skeptic. For instance, in claiming that my belief that I have hands is sensitive, I betray *my conviction* that I'm not a BIV [brain in a vat], either in the actual world or in any nearby worlds. Indeed, I'm ready to admit to the skeptic if I am a BIV, then I don't know I have hands, according to any standards for knowledge. But, of course, *as I firmly believe*, I'm not a BIV (50, emphasis mine).

DeRose claims that the battle over skepticism is about who can provide "the best resolution of our puzzling conflict of intuitions" (49). According to the contextualist, we have one intuition that skeptical arguments are true and valid, and another that we have knowledge in mundane cases. Contextualists aim to provide a theory of knowledge that can account for both these intuitions. In the passage above, DeRose is claiming that we have further intuitions about what is true and what might have been true. Since contextualism aims to provide an account of our intuitions, DeRose thinks he is justified in relying on intuitions about actual and counterfactual worlds as well as our intuitions about when we have knowledge. Thus, DeRose uses our intuitions about actual and counterfactual worlds to construct the possible worlds model referred to by the Rule of Sensitivity.

DeRose's use of intuitions to define knowledge differentiates him from other contextualist theorists in an important way. Traditionally, contextualists take our intuitions as evidence that an unidentified factor makes skeptical hypotheses too remote to be relevant in most contexts. Identifying that factor is the same as explaining the reasons for our intuitions. DeRose rests his definition of knowledge on our intuitions themselves, and thus eliminates the need for an explanatory factor. Put another way, DeRose accounts for our intuitions about knowledge by making those intuitions part of his definition of knowledge.

If DeRose wants intuitions to determine the possible worlds model then he must indicate how intuitions determine unique possible worlds models. DeRose does not supply sufficient textual evidence to indicate how he intends us to do so. Therefore, it is necessary to reconstruct this missing component of his theory.⁶

There are two possible methods for determining a possible worlds model with intuitions. First, one can treat all of a person's intuitions as a set of facts that constitute the actual world, and define distance from the actual world to a possible world as the number of facts both worlds have in common.⁷ For instance, a possible world closest to the actual world is one with only one fact not in common with the actual world. This world is closer than a possible world that shares all but two facts.

One problem with this approach is that it requires one to individuate all of one's intuitions. There is no fixed way to do so. Suppose I have an intuition that my friend Xavier is in Williamstown, and another that Xavier has brown hair. Do I then have an intuition that a brown haired man is in Williamstown? Do I have an intuition that Williamstown is unlike all other towns because Williamstown is the only town with Xavier in it? There are numerous ways to count one's intuitions, and none appears to be preferable. Each way one enumerates one's intuition completely alters the set of intuitions, thereby altering the possible worlds model, and thereby changing what one knows.⁸ DeRose certainly would not want to relativize knowledge to both a context and some as-yet-to-be-determined method for individuating intuitions.

Another problem is that, by simply treating intuitions as elements of a set, and then determining distance based upon the intersection of sets, DeRose would have no way to express the varying strengths of our intuitions. Suppose a set of intuitions contains a subset of strong intuitions, then a world where a few strong intuitions were not satisfied might be closer than a world in which the strong intuitions were satisfied but trivial intuitions were not satisfied. For instance, consider the world where everything is identical except my mail comes at ten in the morning, instead of after lunch, and by a different carrier. This world might be farther than a world in which there is a greatest prime number.⁹

Even if DeRose could provide a method for determining distance that could both individuate our intuitions and weigh stronger intuitions, it is beyond our capacity to make such a complex series of Bayesian approximations in everyday life. This implies that we cannot determine a possible worlds model while engaged in our practices. This implies that we cannot meaningfully talk about knowledge because the meaning of a particular knowledge ascription is relative to an intuitively constructed possible worlds model. When ascribing knowledge we, quite literally, do not know what we were talking about. Consequently, there are rarely knowledge ascriptions that are non-accidental and true. This violates the motivating contextualist insight that we regularly make deliberate

⁶ Another important and unanswered issue for DeRose is *whose* intuitions determines the possible worlds model. It is either the subject's intuitions, or the ascriber's intuitions. On the face of it, both alternatives are equally attractive. In my later objections, I will use cases of first-person attribution so that it is irrelevant whether DeRose intends us to use the subject or attributer's intuitions.

⁷ This possibility arose from a reading of Lewis (1981). It should be stressed again that the reason DeRose must rely upon intuitions about the actual world, as opposed to facts, is that relying on facts begs the question.

⁸ See Lewis (1981) pages 223-224 for a discussion of a related problem.

⁹ Another way to say the same thing is that this method of basing a possible worlds model on intuition ignores the interconnectedness of our intuitions, and instead treats them as discrete entities.

and true knowledge ascriptions. Therefore, this method is unacceptable for basing a possible worlds model on intuitions.

A second possibility is that one can rely exclusively on an intuition about what would most likely have happened if not-P were true. For this method, what are important are intuitions about what might have happened. For instance, when DeRose claims to know he has two hands the only relevant intuition is the intuition that if he did not have hands he would most likely have lost them in a garbage truck accident. This approach avoids the problems mentioned above. It is important to note that what are important are intuitions about likelihood, not intuitions about relevance.¹⁰ What makes the world in which DeRose lost his hands in the garbage truck accident a close world is that his intuitions tell him that this would have been the most likely scenario for him not to have hands.

After we calculate the distance each not-P possible world is from the actual world we still have to decide when a world is sufficiently close to the first not-P world so that it is relevant. This problem can be phrased as follows. Let A be the closest not-P world. Let B be another not-P world. How are we to determine if B is *sufficiently close* to A so that it is relevant? DeRose cannot demand evidence to determine sufficient closeness, since this begs the question in the same way discussed above.

Consider the following example from Harman (1973). My friend Smith tells me that my friend Jones is going to New York City. I believe that Jones is in New York City, and it is true that he is in New York City. In the closest possible world where he is not in New York City I believe that my friend Smith lied to me, but this is only slightly more likely than the possibility that Jones got in a car accident on the way. So in this case the two counterfactuals are sufficiently close so that they are both relevant or irrelevant together. But what if I am slightly distrustful of my friends and I think that the possibility that Smith is lying is significantly more likely? Are the two counterfactuals still sufficiently close? How can we tell? DeRose gives no indication of how ascribers of knowledge are to make this distinction. Since he cannot rely upon evidence, it therefore appears that he again needs to rely upon intuitions.

Based on the preceding analysis we are now in a position to state DeRose's concept of knowledge.

- S knows P iff:
- (1) P is true, and
 - (2) S believes P and would believe not-P for those not-P worlds that are closest, where closeness is determined by intuitions about likelihood, and
 - (3) S's belief P is true in all conversationally relevant worlds. A world is conversationally relevant if its distance from the actual world is less than or equal to the distance determined to be relevant by some conversational standard.

¹⁰ If DeRose were to rely upon intuitions about relevance, instead of intuitions of likelihood, then his theory would amount to little more than a stipulation that knowledge requires true belief, because whenever S claims to know P then necessarily S does not think that any possibilities are relevant other than the ones S has evidence to eliminate. Thus, in all cases where S has a true belief, and claims to know, then necessarily S would know. This is clearly unacceptable, and violates a number of our intuitions about knowledge, including those that DeRose wishes to explain.

The way DeRose sets this up, either (2) or (3) subsumes the other. Either the conversational standards for knowledge will extend beyond the first not-P world, or the first not-P world is further from actuality than the conversation deems relevant. In either case, the more demanding criterion for knowledge takes precedence.

We are now able to identify the flaws in DeRose's approach. Each of these objections stems from DeRose's reliance upon intuitions in (1) and (2). Since DeRose models his Rule of Sensitivity on the sensitivity requirement in Nozick (1981), this might lead one to wonder why Nozick (1981) does not need to rely on intuitions. Nozick is an externalist, and therefore, for Nozick, we do not have to be aware of what we know. DeRose, on the other hand, is trying to verify our intuition that we know in mundane cases, and thus DeRose needs to make it possible for us to be cognizant of when we know. Since DeRose includes a sensitivity requirement, this entails that DeRose must make counterfactuals accessible to us. Demanding evidence about counterfactuals begs the question, so he consequently relies upon intuitions.¹¹ This reliance derails the contextualist project by producing highly counterintuitive results. Since DeRose aims to explain our intuitions about knowledge, the following objections betray fundamental weaknesses in his project.

The first objection is that intuitions about counterfactual worlds are not always determinable. Suppose that S believes P, and that P is true. To satisfy the Rule of Sensitivity, S would have to believe not-P in the closest not-P world. According to the preceding analysis, some not-P worlds are closer than others because we have intuitions about likelihood, and likelihood determines closeness. This means that, of all the possible scenarios where not-P, some scenarios are more likely.

This makes sense once one realizes that if P is true in this world, then what makes not-P worlds different from this world is the presence of a factor not found in this world. DeRose wants to use our intuitions to sort through all possible factors, and rank them in order of likelihood. Let A be a scenario of not-P. Therefore, for S to know P then S's belief that P must at least be sensitive to the most likely A-worlds. The problem is that often we do not have clear intuitions about which As are more likely. For instance, let P be the claim that Bill Clinton is President. What is the most likely scenario A that would result in not-P? Did he die of a heart attack since the last time you checked the news? Did the Senate find him guilty of impeachment charges two years ago? Did he lose the election to Bob Dole? Has there been a military coup without the knowledge of the media? Has he been assassinated and replaced with an imposter? One's intuitions about these possible worlds might not be clear enough that one can select the most likely not-P world. However, the Rule of Sensitivity requires that one's belief be sensitive to the closest worlds in which not-P. Therefore, if one is to know that Bill Clinton is the president one has to be aware of which world is most likely. If one cannot decide which counterfactual worlds are closest then one does not know that Bill Clinton is president. Since people generally do not have intuitions about worlds in which Bill Clinton is not president, it follows that people regularly do not know who the president is.

¹¹ As noted in footnote seven, each objection draws from cases of first person attribution. Thus, each objection will hold regardless of whether DeRose intends to determine the possible worlds model with the subject or attributer's intuitions.

The indeterminacy becomes greater the more certain we are that P. For instance, one probably has very ill defined intuitions about the first world where the earth's gravitation force is twice as strong. DeRose thus ends up having to deny knowledge in the very cases we would most like to grant it.

Moreover, it is integral to DeRose's system that the scope of relevant worlds can expand. But in order for the scope of relevant worlds to expand, one must have already placed possible worlds into order. For instance, suppose that during a conversation about one's belief P an unlikely A is mentioned. To expand the scope of relevant worlds one must be able to determine whether every A-world is more likely or less likely than the mentioned A-world. Since any A-world could be mentioned, one must be able to order the entire set of A-worlds. However, one's intuitions about counterfactuals are likely to be too indeterminate for such an ordering.

Consider the possibility that Clinton has died of a heart attack since you last checked the news. Is this more or less likely than the possibility that he lost the election to Bob Dole? When one's intuitions are indeterminate, one cannot satisfy the Rule of Sensitivity. Thus, DeRose's methodology often forces us to deny knowledge when we intuitively want to grant it.

Second, the use of intuitions about likelihood to determine the closeness of counterfactual worlds makes irrelevant counterfactual worlds relevant. Note that one's selection of A is in no way temporally constrained. The not-P world could have diverged from our own a long time ago, as did the world in which Bob Dole won the election and DeRose lost his hands while working on his uncle's garbage truck.¹²

Recall the case where Smith tells you that your friend Jones has driven to New York City. Suppose that Jones decided to travel to New York City to pick up his reward after winning the lottery. You believe that Jones won the lottery because you saw the ticket and the winning numbers. Even if the odds that Smith is lying are significant, the odds that he would have lost the lottery are greater. Then, in the closest world where Jones is not in New York City, Jones did not win the lottery. Since you would have noticed if he did not win the lottery, you therefore know he is in New York. This is counterintuitive because one would have thought that whether or not you know Jones is in New York would depend upon facts that are more recent.

In this way, events in the past that we usually take for granted become the relevant alternatives on which knowledge claims hinge. The Rule of Sensitivity is often satisfied without a demonstration of sensitivity to what are normally considered relevant counterfactuals. We would like to say that some not-P worlds are irrelevant even if they were more likely. But in order for DeRose to accept this intuition about relevancy, he

¹² Some accounts of counterfactuals allow one to fix closeness from the time of the antecedent. For instance, for the claim "If Nixon had pressed the red button then nuclear holocaust would have followed", one determines the truth-value by looking at all worlds where the antecedent is true, and then determining the closeness of consequents by their similarity to the world at the time of the antecedent, assuming that the laws of physics remain constant. If the consequent is true in all close worlds where the antecedent is also true, then the counterfactual claim is true. DeRose's Rule of Sensitivity, on the other hand, depends upon our holding fixed not-P and identifying the most likely antecedent. The antecedent could have taken place at any number of times, as we see in the worlds where Bill Clinton is not president. There is thus no way to fix temporally the antecedent, and this opens him up to the problem that some scenarios that led to not-P worlds occurred quite a long time ago. Such worlds would normally be assumed to be irrelevant in the course of a knowledge ascription.

needs to find a way to determine which not-P worlds are relevant based upon something besides our intuitions about likelihood.

A third objection is that DeRose makes it cumbersome to talk about knowledge ascriptions. DeRose makes knowledge ascription relative to a possible worlds model. This is central to his claimed ability to solve the skeptical problem. DeRose claims that in most contexts one's true belief only has to be sensitive to near counterfactual worlds in order to qualify as knowledge. In those contexts, what one *means* when one refers to knowledge is that one's true belief is sensitive in close worlds. Further, what counts as a close world is relative to a possible worlds model. Therefore, the meaning of a claim to know is relative to the selection of a possible worlds model. Thus, for two people to talk meaningfully about a knowledge ascription both must be aware of the possible worlds model involved. But, as shown in the first objection, this requires that both people are aware of all the relevant intuitions. This is unlikely to occur. Suppose I say to you, "I know that Bill Clinton is president". You will not know what I mean by this until I have also told you about the strength of all my intuitions about counterfactuals in which Bill Clinton is not president. This could take a long time. In my entire life I have never provided someone else with that kind of information or had another provide that information to me. This means that I have never had a meaningful conversation about what I knew. DeRose's position opens the door to a new kind of skepticism. Traditional skepticism threatened to show that all mundane knowledge ascriptions are false. This new kind of skepticism can successfully show that all mundane knowledge ascriptions are meaningless for everyone except the speaker. This does not seem to be much of an improvement.

Fourth, one's intuitions about counterfactual worlds are not stable. One's intuitions often change depending upon when they are examined. Suppose that S's friend Smith reports that Jones drove to New York City the previous evening. S subsequently forms the belief that Jones is in New York City. Suppose that Jones really is in New York City. S's intuition is that in the closest world where Jones is not in New York City, Jones had a car accident on the way. S's intuition is that if Jones had a car accident then Jones' family would have called S immediately. S's belief is therefore sensitive in the closest world where Jones is not in New York City. S does not need to reflect on this intuition since DeRose does not require a putative knower to be aware that the belief in question is sensitive.¹³ Assuming the standards for knowledge have not been raised, S would therefore know that Jones is in New York City. The next day S wakes up in a terrible mood and distrusts his friends. Perhaps S had a dream in which his friends betrayed him. S never reflects upon his belief that Jones is in New York. However, if S were to reflect on his intuitions then S would claim that in the most likely world where Jones is not in New York City, Smith had lied to him. S's belief is not sensitive in that world since if Smith told a lie, then S would believe it. S therefore no longer knows that Jones is in New York City. S does not have to be conscious of the change in his psychological profile in order to lose the knowledge he had the night before. This is a problem for DeRose's Rule of Sensitivity because, as soon as one stops reflecting on one's intuitions, one cannot be sure that they have not changed and that what one knows subsequently changed. Of course, one might still know P but that is something one cannot be aware of without stopping to reflect once again on not-P worlds. This might be fine for a pure

¹³ Otherwise, since very few people are aware of DeRose's Rule of Sensitivity, and therefore very few reflect on relevantly close counterfactuals, very few people would know.

externalist, but if DeRose wants to retain our intuitions about knowledge in mundane cases then he has to retain our ability to aware of what we know. Otherwise, our routine self-ascriptions of knowledge are non-accidental and true only for the limited period of attribution.

Fifth, intuitions about when two worlds are sufficiently close are seldom well-defined. Recall that earlier it was shown that DeRose needs intuitions not only to determine which counterfactuals worlds are closest, but he needs intuitions to determine which possible worlds are sufficiently close so that they are relevant. Put another way, intuitions not only order the set of all possible worlds but they also divide that ordering into a subset of sufficiently close worlds. The difficulty is that this requires a second set of intuitions. Consider the case discussed earlier about S's belief that Jones is in New York City. Suppose that S's intuition is that it is more likely Jones got into a car crash than that Smith was lying. S then needs a second kind of intuition to determine whether the possibility that Smith was lying is sufficiently close to the possibility that Jones got into a car accident. These possibilities might be too close to call or so distant that the line is easy to draw. However, given the number of possible counterfactuals it is more likely that differentiating between sufficiently close worlds and insufficiently close worlds requires precision. In most cases, one's intuitions cannot supply this level of precision. But if we don't have a precise intuition about what is sufficiently close then, in those cases, the relevance of some counterfactuals cannot be determined. In these cases it would be impossible to make a distinction in principal. But a satisfactory theory of knowledge cannot allow for this kind of indeterminacy since the worlds we can neither exclude nor include might include the very world that determines if the subject knows. DeRose's only option is to deny knowledge to those with indeterminate intuitions about sufficient closeness, but this would require him to deny knowledge to a majority of the mundane cases. DeRose would have to deny knowledge to those that have good evidence for a belief, but whose intuitions about sufficient closeness are indeterminate. Given the uncertain nature of intuitions about sufficient closeness, this requires DeRose to deny knowledge in most cases.

Finally, the decisive role of intuitions in The Rule of Sensitivity requires one to ascribe knowledge to those without good evidence. Suppose that I have a very particular belief about the trajectory of a meteor based on astrological charts. Suppose that I have strong faith in the power of astrological charts, and in my ability to interpret them. Then in the closest world where the meteor does not have the trajectory I predict, I believe my astrological charts would reflect this fact and I would have formed a different belief about the meteor's trajectory. Therefore, my belief is sensitive. Suppose that my belief is miraculously true. I therefore know the trajectory of the meteor. This is a counterintuitive result. It quickly leads to a kind of epistemological relativism where knowledge claims are entirely dependent upon who is making them. The reason that knowledge is traditionally thought to have a justification requirement that demands evidence is that we must avoid this type of problem. In DeRose's analysis, if the possible worlds model is bizarre, but a subject's belief is true, then they could know nonetheless.

On the other hand, suppose I believe no greatest prime number exists. I have memorized the proof and can recite it whenever someone asks about it. In the closest world where there is a greatest prime number, I believe that the rules of mathematics must have altered. But my intuition is that such a world must be so distant and bizarre that I might still believe that there is no greatest prime number. My belief therefore does

not satisfy the Rule of Sensitivity, and therefore I do not know that there is no greatest prime number. Likewise consider the Leibnizian that believes God chose this world over all other possible worlds because this world is perfect. For the Leibnizian, if things had turned out differently then God would have made an error of judgment, and in any world where God is capable of making errors in judgment there can be no guarantee that what beliefs will track the truth. Consequently, the Leibnizian's beliefs can never satisfy the Rule of Sensitivity. Therefore a Leibnizian can never truthfully ascribe knowledge to herself, regardless of how good her evidence is. This is counterintuitive.

DeRose's Rule of Sensitivity precludes us from knowing some of the facts about which we are the most confident, and allows others to grant knowledge to themselves when we would like to deny them that ability. Both of these lapses are caused by DeRose's excessive reliance on intuitions, and his subsequent neglect of the traditional evidential requirement.

These objections give us reason to reject DeRose's contextualism and any other contextualist strategy that uses intuitions to determine a possible worlds model acted upon by a sensitivity requirement. One can easily go back through each objection and replace "intuition" with "belief" without affecting the arguments. Therefore, one cannot develop a contextualist strategy that uses beliefs instead of intuitions to determine a possible worlds model. The contextualist insistence that we regularly make true knowledge claims prevents them from using facts, and their insistence that the skeptic is sometimes successful prevents them from using evidence. The only way the contextualist could retain the sensitivity requirement would be if she could identify something to determine the possible worlds model besides evidence, beliefs, intuitions, or facts. It is unclear what this leaves the contextualist to work with.

DeRose's failure demonstrates that the contextualist cannot retain Nozick's sensitivity requirement while maintaining that her theory validates our intuitions about knowledge.

Stalnaker and Pragmatic Presuppositions

In the previous section I showed that contextualists cannot rely on a Nozickean sensitivity requirement to explain variances in the meaning of knowledge ascriptions. Contextualists therefore need to provide an alternative definition for the “context” of a knowledge ascription. Stalnaker (1972) and (1973) define contexts in terms of pragmatic presuppositions, thereby supplying an alternative framework within which one might develop a contextualist theory of knowledge. In this section I argue that a contextualist theory of knowledge built upon Stalnaker’s definition of context does not withstand scrutiny.

It is important to discuss Stalnaker because several prominent contextualists have suggested that his work is a way of completing a contextualist theory of knowledge. Unfortunately, deriving a contextualist theory of knowledge from Stalnaker’s work requires a significant amount of reconstruction. Therefore, it is difficult to be sure whether Stalnaker would accept my considerable elaborations. Assume that all passages without citation are the product of my own reflections on Stalnaker.

In the course of any conversation a great deal is presupposed. Consider the sentence “Jim’s dogs are asleep, and Jim has dogs”.¹ The phrase “Jim has dogs” is ordinarily redundant since this fact was already presumed by the statement, “Jim’s dogs are asleep”. Likewise, suppose that I say, “Do you think Jim will win the competition?” and you respond, “No, Bob has been running laps all year”. This response does not make sense unless we both believe that Bob is in the race, that Bob has a chance to beat Jim, and that one’s chances in the race are improved by running laps.

Stalnaker (1973) differentiates two types of presuppositions: semantic presuppositions and pragmatic presuppositions. A presupposition is semantic if it can be explained in terms of the content of sentences. The fact that Jim has dogs is a semantic presupposition of “Jim’s dogs are asleep.” A presupposition is pragmatic if it must be explained partly in terms of facts about the beliefs, intuitions, and expectations of users of sentences (447).² The response “Bob has been running laps all year” depends upon pragmatic presuppositions.

Stalnaker writes about pragmatic presuppositions:

Statements and requests are made, questions asked, proclamations and commands issued, against a background of common knowledge, or at least what is represented as common knowledge. The background of knowledge or beliefs purportedly shared by the speaker and his audience constitute the presuppositions which define context. A rough definition might go something like this: *A speaker presupposes that P at a given moment in a conversation just in case he is*

¹ A similar example is found in Lewis (1979)

² Stalnaker (1973) points out that, strictly speaking, semantic presuppositions might be a subset of pragmatic presuppositions. However, it is useful distinction to maintain.

disposed to act, in his linguistic behavior, as if he takes the truth of P for granted, and as if he assumes that his audience recognizes that he is doing so (448, emphasis in original).

Stalnaker's account of pragmatic presupposition has two important features. First, it is sufficient that speakers in a conversation are *disposed* to presuppose certain propositions. Stalnaker writes, "Because I *would* act, in my linguistic behavior, as if I take the truth of that proposition for granted in these ways, I *do* in fact presuppose it" (449, emphasis in original). Second, context is defined as what the speaker takes to be the common presuppositions of participants in a conversation. "Conversation" is defined broadly enough so that it includes conversations with oneself and conversations between authors and readers.

According to Stalnaker, the fiction of possible worlds is a powerful apparatus for unifying philosophical analysis (see 1984, 1-25). For instance, in *Inquiry* Stalnaker defines belief, desire, truth, proposition, and intentionality in terms of possible worlds, thereby giving him a useful tool for interrelating the five concepts. Moreover, Stalnaker argues that formalizing philosophical inquiry in the terms of possible worlds analysis only requires a commitment to a few, intuitively acceptable axioms (see 1984, 43-58). If one accepts both claims, then Stalnaker has provided grounds for accepting possible worlds as a foundational philosophical concept.

Pragmatic presupposition might be thought about in terms of possible worlds. If a possible world is consistent with all presuppositions then it is relevant in the conversation. If it is inconsistent with any presupposition then it is irrelevant (Stalnaker 1973, 450). If one views conversations as rule-governed activities, then a proposition is permissible if and only if it is a proposition about a relevant possible world. Thus, as the set of presuppositions changes, there is a corresponding shift in the set of relevant worlds. This in turn alters the set of permissible propositions.

In order to develop an understanding of how one might try to use Stalnaker's analysis of pragmatic presuppositions as the starting point for a contextualist theory of knowledge, it is useful to consider how Stalnaker defines a proposition in *Inquiry*. He writes, "a proposition is a function from possible worlds into truth-values" (1984, 2). In other words, a proposition is nothing more than a rule for selecting a subset from the set of all possible worlds.³ Further, the meaning of a proposition corresponds to the function, or rule, associated with a particular proposition. Thus, if two propositions have identical functions from possible worlds to truth-values, then the meaning of the two propositions are identical. For instance, "the Morning Star is rising over the

³ If one finds the language of possible worlds, sets, and functions mystifying, then it might be helpful to think of propositions as entities with truth-conditions, and associate the meaning of a proposition with its particular truth-conditions. This is a distortion of Stalnaker's definition, but it is a familiar definition, and for our purposes it will not affect the arguments. However, if one chooses to think of propositions as entities with truth-conditions, instead of functions from possible worlds to truth-values, then one should keep in mind what seem to me to be two reasons why Stalnaker's actual definition is superior. First, defining a proposition as an entity with truth-conditions is circular since truth-conditions are necessarily in the form of propositions. Second, this definition of propositions places a non-trivial importance on the term "truth", and "truth" has proven to be a difficult concept to treat formally. Stalnaker's avoids the first problem by defining propositions in terms of possible worlds. He avoids the second problem by making truth-values nothing more than the values in the range of the function. In other words, a proposition divides possible worlds into two sets, and we arbitrarily assign all worlds in one set "true" and the other "false".

horizon” and “the Evening Star is rising over the horizon” will be true in the same set of possible worlds, and therefore both propositions have the same meaning. This definition of propositional meaning has a number of advantages, most notably that it does not expressly rely on a definition of intentionality.⁴

Putting this together with Stalnaker’s definition of pragmatic presuppositions, one might claim that the set of contextually relevant worlds *partially* determines the meaning of a claim about propositional knowledge. I emphasize “partially” because, while propositional meaning is a function from possible worlds to truth-values, pragmatic presupposition is a function from possible worlds to relevance-values (i.e. relevant or irrelevant). Stalnaker claims that the meaning of a sentence often depends upon what is relevant at the time. The meaning of a proposition is therefore often a composition of two functions, one that maps all possible worlds to relevant or irrelevant, and one that maps all relevant worlds to true or false. I will claim later that propositions about knowledge are a composite function of this form.

Think of possible worlds as discrete drops of water in an infinite ocean. Divide the ocean into two sections. One section that contains all drops consistent with every pragmatic presupposition, and a second section containing all other drops. Take the section that contains the drops consistent with every pragmatic presupposition and assign each drop a value of “true” or “false”. Color those that are “true” with black dye, and those that are “false” with a white dye. The meaning of a proposition is the process that assigns colors to particular drops from the infinite ocean. If two processes assign the same colors to the same drops, then, for our purposes, they are the same process. For our purposes, this would be a process of assigning truth-values to contextually relevant worlds. Moreover, if a process places more drops in the “relevant” section then necessarily more drops will have coloring, so the process is different from one that places less drops in the “relevant” section.

Understanding possible worlds as drops of water allows one to think of a function as a physical process of sorting through a set. However, one should be careful not to think about functions to possible worlds as a similar kind of physical sorting, or one might begin to ask questions like, “How do we access each possible world in order to assign it a truth-value?” Functions from possible worlds to truth-values are relations between sets of abstract entities, similar to the way $y=3/5x+5$ relates a set of real numbers to another set of real numbers.⁵

An important implication of this analysis is that presupposing underdetermines the meaning of a proposition. Thinking of this in terms of the ocean metaphor: once one has divided the ocean into a “relevant” section and an “irrelevant” section, there are many different ways one can color “relevant” drops with white or black. Each of these different ways will alter the overall function. Thinking of this in terms of the composition of functions: if $g(x)$ and $f(x)$ are different functions, then $g(h(x))$ does not equal $f(h(x))$.

It is left an open question whether a speaker has the ability to choose what her proposition means. If one thinks that the intentions of the speaker determine what proposition corresponds to her sentence, then one can allow for a plurality of meanings

⁴ According to Stalnaker, those who rest a theory of propositional meaning on an account of intentionality, are unable to turn around and define intentions as mental states that map to propositions. This is problematic because it is unclear how else one could define intentionality.

⁵ Strictly speaking there is only one set of the real numbers, but that is besides the point.

for the same pragmatic presuppositions.⁶ If one thinks the relevant linguistic community determines what proposition corresponds to a speaker's sentence, then once the pragmatic presuppositions are set, she can mean only one proposition. One can also try to negotiate the middle ground between these two options by letting the community fix propositions to sentences for some sentences and not others. It is a useful feature of Stalnaker's analysis that it is noncommittal on this question.

A contextualist theory of knowledge follows straightforwardly from this analysis. If the contextualist can establish that "S knows X" has several potential meanings, then Stalnaker's account of pragmatic presupposition would allow her explain what is responsible for the different meanings. In particular, Stalnaker's account of pragmatic presupposition allows the contextualist to understand changes in the meaning of knowledge claims as changes in pragmatic presuppositions. For example, a skeptic in an epistemology classroom and a grocer asking for directions probably mean very different things when they speak of knowledge. This is due to discrepancies in the set of facts they are willing to take for granted. Thus, when skeptics and non-skeptics get into an argument, in all likelihood they are asserting and denying different propositions.

Stalnaker's framework leaves it an open question whether there is more than one standard for justification. Suppose that P_1 and P_2 are propositions of the form "S knows X" that are uttered with the same set of presuppositions, the same S, and the same X. If there is a possible world that P_1 sends to "true" and P_2 sends to "false", then P_1 and P_2 have different requirements for S to know X. Suppose that we define knowledge as a justified true belief. Then, because the truth and belief condition would be satisfied (or unsatisfied) for both propositions, P_1 and P_2 must have different justification requirements. On the other hand, if P_1 and P_2 necessarily send all worlds to the same truth-values, then justification is an invariant condition. It is important to note that invariantists about justification can still be contextualists about knowledge because they acknowledge the existence of many different, equally acceptable, sets of pragmatic presuppositions.

In his paper "Discrimination and Perceptual Knowledge", Goldman spends one paragraph sketching a contextualist theory of knowledge that builds on Stalnaker's account of pragmatic presupposition (1976, 776-777). In the next paragraph he objects to it on the grounds that, "[I]f a speaker utters a knowledge sentence without presupposing a fully determinate set of alternatives, he does not assert or deny any proposition. That seems too strong" (777). It is difficult to make sense of what Goldman means by "presupposing a fully determinate set of alternatives" because one presupposes facts, not alternatives. Moreover, the set of "alternatives" is just the set of possible worlds consistent with the facts the speaker presupposes. The only way for the "set of alternatives" to be less than "fully determinate" is if the presuppositions are indeterminate. Recall that Stalnaker defines a speaker's presuppositions as follows: "A speaker presupposes that P at a given moment in a conversation just in case he is disposed to act, in his linguistic behavior, as if he takes the truth of P for granted, and as if he assumes that his audience recognizes that he is doing so" (448, originally all in italics). Therefore, if a speaker has indeterminate presuppositions then she must have at least one indeterminate disposition to act as though something is true. Thus, for

⁶ I am thinking here of those attracted to the views of Grice (1989).

Goldman to be correct, there must be cases where the speaker has a disposition, but it is impossible to tell whether it is a disposition to take something as true or false.

There are two types of indeterminacy: epistemological indeterminacy and metaphysical indeterminacy. Contextualists are committed to the claim that we can be aware of what we know. Therefore, for the Goldman's argument to be successful, it is sufficient to show that the indeterminacy in question is epistemological.

If Goldman's argument is that dispositions are epistemologically indeterminate, then he is on the right track. Consider the difference between the following two statements: "S will do X", and "S is disposed to do X". The difference is that if S is disposed to do X then S will not necessarily do X. In other words, there exist conditions under which S will not do X. Consider the statement, "S is disposed to obesity". This implies that it is possible that S will not become obese. Perhaps S could go on a life-long diet.

How can one determine what an attributer of knowledge is disposed to take as true?⁷ One might think that talking with the attributer can reveal what she was disposed to take as true at the time of the attribution. This, one might hope, allows us to determine the meaning of a knowledge ascription retroactively. However, one of the premises of contextualism is that, if we give the skeptic the ability to make her case, people will generally agree that she is right. Thus, people are generally disposed to take the skeptical argument as true. Of course, if you talk to people about baseball or the weather, they are likely to implicitly treat skeptical hypotheses as false. But why should this make a difference? One person can have many conflicting dispositions. For instance, I am disposed to be cordial when smiled at, but disposed to be rude when laughed at. Why should one give the non-skeptic the upper hand by stipulating that only non-skeptical conversations can determine the attributer's dispositions? Therefore, determining dispositions based on future conversations is not a good method.

This might lead one to think that facts about the attributer's past conversations will reveal her dispositions in the present case. However, this runs into an induction problem. Previously it was demonstrated that what differentiates "S is disposed to do X" from "S will do X" is that there are circumstances in which S will not do X. In order to determine what S is disposed to do next in the conversation, one must therefore determine under what circumstances S will do X. The problem is that based on past cases alone it is impossible to tell what those circumstances are. Consider the claim, "Alex is disposed to get angry when drunk." One might think that after watching Alex get drunk enough times in the past, and enduring his subsequent petulance, that one could reasonably make such a claim. However, this is not enough to justify the claim that Alex is disposed to get angry when drunk. Perhaps in the past Alex has only gotten drunk after a relationship ended badly, and if he were to get drunk during less turbulent times he would not become angry. Perhaps Alex gains peace of mind after he graduates from college, so it is misleading to consider cases in college when he became angry while drunk. Perhaps Alex only gets angry when drunk if he runs out of beer, and if one gives him enough beer, he will remain placid. It just so happens that until now, Alex has exhausted the beer supply every time he drank alcohol. As has been shown convincingly by Goodman's (1965) discussion of "grue", past cases always allow for equally justified and contradictory inductive arguments. This undermines one's ability to determine dispositions based on past cases.

⁷ The following analysis explicitly addresses cases where one is not the attributer, but it generalizes to handle cases where one is the attributer.

Further, determining someone's presuppositions based on past cases runs into a regress problem. Stalnaker defines the context of an utterance as all of the speaker's current dispositions to take as true. The meaning of a speaker's utterances depends in part on its context. Therefore, in order to figure out what a speaker meant in past cases, one would have had to figure out what her dispositions were in those cases. However, if figuring out a speaker's dispositions is always a matter of looking at past cases, then one will need to base one's assessment on even older cases. Determining what she meant for those cases will require one to look at yet older cases, and so on. Therefore, past cases will not allow one to determine a speaker's present dispositions.

Goldman's objection, as I have developed it, is therefore correct. A speaker's dispositions are often, if not always, epistemically indeterminate. This makes the meaning of knowledge ascriptions indeterminate. Therefore, a contextualist cannot claim that our knowledge ascriptions are regularly true while building upon Stalnaker's account of pragmatic presuppositions. A contextualist attracted to the way in which Stalnaker explains variances in the meaning of "know" needs to offer an alternate account of presupposing. This is precisely what Lewis (1979) provides. Lewis (1996) attempts to develop a theory of knowledge from within his (1979) framework. In the next section, we will consider both of Lewis' papers.

Lewis' "Elusive Knowledge"

A contextualist argues that the meaning of a knowledge ascription depends on the context. Stalnaker (1973) defines the context of a sentence as the current set of presuppositions. According to Stalnaker, a speaker presupposes whatever she is disposed to take for granted. Section four demonstrated why this is an unsuitable account of presuppositions if one wants to explain variances in the meaning of knowledge ascriptions in terms of presuppositions.

Like Stalnaker, Lewis (1979) defines the context of a sentence as the set of current presuppositions. Unlike Stalnaker, Lewis defines presuppositions in terms of what sentences participants in a conversation have already uttered. A helpful way of thinking about the difference between Stalnaker's and Lewis' approaches to presuppositions is that Stalnaker bases presuppositions on what might be said, while Lewis bases presuppositions on what has been said.

In order to help explain how presuppositions can change, Lewis introduces the concept of the "conversational score". A conversational score is a set-theoretic construct that contains all current presuppositions (345). This includes presuppositions about what is true, what indexicals like "she" presently refer to, and the particular sense given to words like "empty" and "flat".

A speech act is permissible if it is consistent with the conversational score. For instance, returning to an earlier example, suppose that I ask, "Do you think Jim will win the competition?" and you respond, "No, Bob has been running laps all year". Your response is impermissible if it presupposes facts that were not established previously in the conversation. One such presupposition might be the fact that Bob is competing against Jim. For your response to be permissible, someone must introduce that fact into the conversation, which Lewis understands as an addition to the conversational score. One may also add to the conversational score through hypothesis, or subtract from the conversational score by questioning something that was previously taken for granted.

According to Lewis (1979) the conversational score typically will be adjusted so that the most recent utterance made by a participant in the conversation can be true. He calls this principle the Rule of Accommodation. For instance, the word "can" has a relative modality that is dependent upon the context (354-355). If someone says "I see a cat in front of me and therefore I can't be wrong in my belief that there is a cat in front of me" then that person is presupposing a particular modality for "can". If someone responds, "You might be dreaming" then she is attempting to shift the presuppositions behind the word "can". If the shift were successful then the first person would have to admit that her belief, that there is a cat in front of her, could be wrong. But this does not imply that her original claim was false. It was true under the previous set of presuppositions but false under the current set.

Lewis (1996) builds a theory of knowledge on his (1979) account of conversational scores (1996, 515). He defines knowledge as follows: "S knows that P iff S's evidence

eliminates every possibility in which not-p – Psst! – except for those possibilities that conflict with our proper presuppositions” (1996, 506).

According to Lewis, whether or not one knows P often depends upon which not-P possibilities are presupposed as not obtaining, because if a proposition is not presupposed then evidence is needed to eliminate it. Suppose that one is at the zoo with one's son.¹ While looking at a striped animal in a pen labeled “zebras”, one can truthfully say, “I know the animal I am seeing is a zebra” just in case it is presupposed that it is not a painted mule. If someone then asks “But couldn't it be a painted mule?” then this possibility is no longer presupposed. Unless one can then produce evidence that it is not a painted mule, then one must admit that one does not know it is a zebra. This provides an example of how Lewis (1996) deals with skeptical arguments. If a skeptical hypothesis is presupposed then one does not need evidence to eliminate it. One may simply ignore it.

Of course, just because a possibility is presupposed does not necessarily imply that it is *properly* presupposed. For instance, when ascribing knowledge one might accidentally presuppose that something true is false. If it really is a painted mule then, no matter what the conversational presuppositions are, one cannot truthfully say, “I know the animal I am seeing is a zebra”. In the case of knowledge ascription, some presuppositions are improper regardless of the presuppositions of a speaker and her audience.

Lewis (1996) is an attempt to identify rules that can distinguish improper presuppositions from proper presuppositions for knowledge ascription. He identifies six rules for determining what may be properly presupposed when ascribing knowledge. They are as follows:

- | | |
|-------------------------------|---|
| (1) The Rule of Actuality: | “The possibility that actually obtains is never properly ignored” (506). |
| (2) The Rule of Belief: | “A possibility that the subject believes to obtain is not properly ignored” (507). |
| (3) The Rule of Resemblance: | “Suppose that one possibility saliently resembles another. Then if one of them may not be properly ignored, neither may the other” (508). |
| (4) The Rule of Reliability: | We may ignore the possibility that reliable processes are malfunctioning. |
| (5), (6) Rules of Method | “We are entitled to presuppose ... that a sample is representative; and that the best explanation of our evidence is the true explanation” (509). |
| (7) The Rule of Conservatism: | Possibilities normally ignored by those around us may be ignored. |
| (8) The Rule of Attention: | “A possibility not ignored at all is <i>ipso facto</i> not properly ignored” (509). |

¹ This example is introduced into the literature in Dretske (1970), and was described in the introduction. Stine (1971) and Vogel (1990) offer commentaries on Dretske's treatment of the case.

The Rule of Actuality functions as the truth condition preventing one from claiming S knows P when P is false. The Rule of Belief functions as the belief condition preventing us from claiming S knows P when S does not believe P. Both the Rule of Actuality and Rule of Belief conform to the two traditionally uncontroversial criteria for knowledge.

The Rule of Resemblance mandates that if two possibilities saliently resemble each other, and one of them is not properly presupposed, then the other is not properly presupposed. It is a second-order rule, operating on the set of properly ignored possibilities determined by the other seven rules. It is central for Lewis' claimed ability to solve Gettier cases.²

In order to give the Rule of Resemblance a fully determined sense, it is necessary that Lewis provide a stable criterion for salience. Lewis, however, admits that the concept of salience is itself contextually sensitive (514). Moreover, not all resemblances between possibilities can be counted. For instance, the possibility that is true and the possibility of being fooled by an omnipotent demon resemble each other in one respect: neither can be successfully eliminated through a subject's evidence (508). But if such resemblances were counted as salient then skeptical hypotheses could never be properly presupposed.

The next rules are the Rule of Reliability and the Rule of Conservatism. The Rule of Reliability is that one may defeasibly ignore the possibility that reliable processes are malfunctioning. Lewis leaves it vague how we are to generalize each process.³ The likely way for him to solve this problem is to make a selection of the correct process type a function of the context within which the ascription takes place.

The Rule of Conservatism is that "generally ignored possibilities may be properly ignored" (509). Of course, when Lewis refers to "generally" ignored possibilities he is not referring to possibilities that are ignored in *every* context. If that were the case then, because of the skeptical conversations of some epistemologists, the Rule of Conservatism would be vacuous. On the other hand, the concept of possibilities that are "normally ignored" suggests a regularity that cannot be found in the presuppositions in a conversation at a particular moment. The Rule of Conservatism thus also depends upon a prior specification of context.

Both the Rule of Reliability and Rule of Conservatism determine what may be ignored while the first, second, and eighth rules determine what may not be ignored. The Rule of Reliability and Rule of Conservatism are what Lewis labels "permissive rules" (509). In cases of conflict permissive rules always lose. For instance, I can generally trust my vision and therefore by the Rule of Reliability I can presuppose that my vision is reliable. Suppose that I see a cat in my field of vision. Further, suppose that someone slipped a hallucinogen into my drink and that this is what caused me to see the cat. The Rule of Actuality trumps the Rule of Reliability and therefore I cannot ignore the possibility that my vision of the cat is the result of a hallucination (509). It is unclear why Lewis includes permissive rules at all. He could have simply included the observation that a presupposition is a proper presupposition if it violates none of the

² Cohen (1998) argues that Lewis' Rule of Resemblance is unable to solve Gettier cases (527). If one accepts Cohen's objections then one has even more reason for agreeing with me that the Rule of Relevance is contextually sensitive, since only by making the context do much of the work could Lewis potentially answer Cohen's objections.

³ The generality problem for reliabilism is outlined in Conee and Feldman (1998).

non-permissive rules. Including rules that defeasibly validate conversational presuppositions adds nothing substantive to his account, except perhaps some elucidation as to how we can expect normal conversations to run.

As Lewis acknowledges, the Rules of Method can be understood as variants on the Rule of Reliability (509). Therefore, from this point onward, the Rules of Method will be treated as a special case of the Rule of Reliability.

The Rule of Attention is that if a possibility is not ignored then it is not properly presupposed. Thus, whatever participants in a conversation do not presuppose they necessarily do not properly presuppose. To put this in terms of a conversational score, if a presupposition is not an element of the conversational score then necessarily it is not a proper element of the conversational score. This should not be confused with the converse or the inverse, both of which Lewis denies. There is one important application of the Rule of Attention that deserves notice. Suppose one is well informed about S's strange intuitions about counterfactuals. Then, necessarily, when ascribing knowledge to S one's attention focuses on those counterfactuals. Therefore, the Rule of Attention does not allow one to presuppose properly those counterfactuals when ascribing knowledge to S (511).

Traditionally, contextualists thought of the standards for knowledge as being either high or low. This required them to provide a mechanism to rank all competing alternatives so that in cases when the standards for knowledge were low only a few would be relevant, while in cases where the standards for knowledge were high less probable alternatives would become relevant. For instance, Stine (1976) asserts that evidence can be used to order all possibilities, while DeRose (1995) implies that intuitions can serve the same purpose. For both Stine and DeRose, the context determines where to draw the line separating relevant from irrelevant possibilities. In this respect, Lewis (1996) differs greatly from the contextualist tradition. In Lewis' account not all standards are necessarily orderable on a linear scale, where higher standards for knowledge necessarily contain all of the requirements of lower standards. Lewis' rules make it possible for different types of doubts to be important in different contexts because his rules do not rely upon an ordering of possibilities.

Before considering Lewis' proposal in more detail, it is important to understand why Lewis (1996) is immune to the objections raised against DeRose (1995) and the problem identified with a contextualist theory based on Stalnaker's account of presuppositions. First, as mentioned above, Lewis does not require an ordering of the set of relevant possible worlds. According to Lewis, if S knows P then S is required to have evidence to eliminate all possibilities that are not properly presupposed. That is, S is required to have evidence that eliminates all relevant not-P worlds. On this account, a possibility is either properly presupposed or improperly presupposed; correspondingly a possible world is either relevant or irrelevant. There is no gradation of relevance. There is thus no need to select one possible world as actual, and no need to rank all other worlds accordingly. While two worlds can have a fixed distance from each other, perhaps by measuring the number of facts they have in common, there is no need to compare all worlds in a uniform way.⁴ For DeRose, however, the distance of a possible world is measured with

⁴ Lewis thinks that there is an actual world and that possible worlds do have a distance from the actual world. This is central to his analysis of counterfactual conditionals (see Lewis 1986). However, his account of knowledge does not require that we are aware of this ordering.

respect to the actual world. This is what enables him to order all possible worlds and thereby gives sense to concepts like "closest world" and "distant world". It also requires him to rely upon intuitions since he grants that evidence cannot be used to disprove skeptical hypotheses. This causes him difficulty because intuitions are too indeterminate and unreliable to order the set of all possible worlds, and then determine which are sufficiently close. Intuitions are too capricious and irrational to offer the kind of stability that our routine knowledge ascriptions presuppose. By making ordering superfluous, Lewis avoids these problems.

Second, Lewis determines the set of presumptions based on facts about the conversation between an ascriber and her audience. Therefore, if two people are conversing about a knowledge ascription then necessarily they will share the same set of relevant possible worlds. DeRose and Stalnaker, on the other hand, relativize knowledge to a set of private intuitions or dispositions. At best, this makes the meaning of a knowledge ascription cumbersome to communicate. At worst, it makes it thoroughly inaccessible. By using conversations instead of intuitions to define the set of relevant possible worlds, Lewis avoids this trouble.

Third, Lewis does not require a ranking of all counterfactuals by likelihood. Recall that one of the problems facing DeRose was that for the conditional "A therefore not-P" there are no temporal constraints on the selection of A. Therefore, some event could have happened a long time ago, and be taken for granted in the current conversation, but nonetheless on DeRose's model would be the first not-P world and thus potentially the only relevant world for determining sensitivity. Lewis' account of knowledge, however, does not depend upon intuitions about likelihood. Instead, it defeasibly rests upon the conversational presuppositions. Therefore, if during our conversation we ignore some possibility, and it does not conflict with the first three rules, then we have properly presupposed it. For instance, if Jones had won the lottery but Smith and I take that for granted, then the worlds in which Jones lost the lottery are not relevant counterfactual worlds.

Fourth, Lewis includes an evidence requirement for knowledge. For Lewis, once a set of proper presuppositions is determined, then whether or not S knows P will depend on whether S has sufficient evidence to rule out all relevant not-P worlds. DeRose, on the other hand, substitutes a sensitivity requirement for the traditional justification requirement, and consequently allows S to know P even if S has no good evidence that P. Besides being *prima facie* counterintuitive, it also leads to a number of problems when combined with DeRose's method of determining possible worlds models. Lewis requires evidence for knowledge.

Unfortunately, Lewis' analysis runs into a different problem. The Rule of Resemblance, the Rule of Reliability, the Rules of Method, the Rule of Conservatism, and the Rule of Attention all vary across contexts. Building upon our earlier analysis of these rules, this implies that the following concepts are contextually determined: salience, ignoring, a reliable process, and normal behavior. Lewis defines the context of a sentence as the presuppositions at the time someone utters it. Therefore, these concepts can vary with the presuppositions of a conversation. In general, the Rule of Accommodation shifts the set of presuppositions so that the most recent utterance can be true. Knowledge ascription is a kind of utterance, and thus, in accordance with the Rule of Accommodation, the presuppositions will generally shift so that a knowledge ascription can be true. In other words, whenever someone ascribes knowledge, the

conversational presuppositions will shift so the ascription can be true. In the remainder of this section I will demonstrate how the concepts of salience, ignoring, a reliable process, and normal behavior can always be adjusted through presuppositions in such a way that all of Lewis' rules, except the Rule of Actuality and Rule of Belief, can always be satisfied. Therefore, if the Rule of Accommodation always shifts the presuppositions governing salience, ignoring, etc. then the Rule of Resemblance, Rule of Reliability, Rule of Conservatism, and Rule of Attention are vacuous. I will conclude that if Lewis (1996) relies upon Lewis (1979) for an account of presupposing and context, then all that Lewis (1996) requires for knowledge is a true belief and a knowledge ascription.

Consider Ginet's barn case, introduced into the literature by Goldman (1976). Suppose Henry is driving through a district where the majority of what appear to be barns are papier-mâché facsimiles of barns. Henry looks outside his window and sees one of the few genuine barns, and subsequently forms the true belief that he is seeing a barn. However, if it had been a papier-mâché facsimile of a barn then he still would have believed it was a barn. Goldman argues that Henry does not know he is seeing a barn because his belief is accidentally true.

Lewis claims that a combination of the Rule of Actuality and Rule of Resemblance can handle this case because the possibility that it is a bogus barn "resembles actuality in respect of the abundance of bogus barns, and the scarcity of real ones, hereabouts" (508). Why does the possibility that one is seeing a papier-mâché barn resemble the possibility that one is seeing a real barn? Because, as long as one is in that district, the *probability* of seeing a papier-mâché barn is quite high, thereby making this possibility relevant.⁵ Lewis thinks that he can treat Gettier cases and the lottery case similarly because in each case, the probability that one's belief P could be false is high enough that this not-P possibility saliently resembles actuality.

However, if conversational presuppositions shift so that Henry's knowledge claim can be true, then the presuppositions about "salience" will shift so that what he claims can be true. Perhaps the standards for salience will shift so that the probability of seeing a papier-mâché barn is calculated by considering the rate of papier-mâché barns nationwide. Or, if there is no barn facsimile in sight, the probability can be calculated based on the chances of seeing a papier-mâché barn within the region of Henry's sight. Thus, even when in papier-mâché barn country, the possibility that Henry is seeing a papier-mâché barn does not necessarily saliently resemble the possibility that Henry is seeing a real barn.⁶ It only becomes relevant if the attributer determines the probability based on the set of papier-mâché barn cases in the district. If Henry ascribes knowledge to himself then the presuppositions will shift so that the probability of his seeing a papier-mâché barn is not based upon the frequency of those barns within the papier-mâché district. The presuppositions will shift so that what he said can be true. In a similar way, the presuppositions regarding salience can always shift so that whatever knowledge claim is made does not violate the Rule of Resemblance.

Similarly, the Rule of Conservatism can shift so that it accords with Henry's claim to know he is seeing a real barn. The Rule of Conservatism allows one to ignore possibilities that those around us normally ignore. Suppose that the residents of papier-mâché barn county do not ignore the possibility that what appears to be a barn might be made of papier-mâché. But the phrase "normally ignore" is vague enough so that the

⁵ In this regard, Ginet's barn case is equivalent to Dretske's (1981) Gadwall Duck example.

⁶ This argument came from a reading of Brandom's (2000), and Conee and Feldman's (1998).

presuppositions fixing it can be shifted so that we consider what everyone nationwide normally ignores, or what people in Henry's car normally ignore. Thereby, the habits of the resident of papier-mâché barn country would become either insignificant or irrelevant.

The problem is that the phrase "Henry seeing the barn" can be identified as a token of any of the following types "All people seeing that object", "Some people seeing that object", "All people seeing very similar objects", "Some people seeing very similar objects", "All people seeing somewhat similar objects", or "Some people seeing somewhat similar objects". Each of these tokens will produce a different set of normal behavior, and thereby make it defeasibly permissible to ignore a different set of possibilities. By the Rule of Accommodation, the conversational presuppositions shifts so that whatever token validates the last utterance will become part of the conversational score. Thus, the Rule of Conservatism is moot unless no community of people can be defined that ignores the same things as the putative knower. Given the flexibility of the Rule of Conservatism, and the possibility of defining a community including just the putative knower and her past ignorings, the only time this will happen is when the fact the putative knower claims to know is false. But we already have the Rule of Actuality to take care of such cases.

Similarly, the Rule of Reliability allows us to ignore the possibility that reliable processes are malfunctioning. It is left up to the context to fix what is to count as a reliable process. If I claim to know I am seeing a barn while in papier-mâché barn district, then the presuppositions change so that I am relying on the general process of seeing. If someone else in the car were subsequently to deny me knowledge they would shift the presuppositions so that my seeing the barn was classified as a token of the process of seeing barns from a moving vehicle in papier-mâché barn country.

The Rule of Attention is the most straightforwardly contextual. Lewis claims that what counts as ignoring is relative to the context. Thus, the standards for ignoring will always shift so that someone who ascribes knowledge does not neglect to have evidence to eliminate an ignored alternative. Suppose that someone in the car mentions that this is a region replete with papier-mâché barns. Then, a few moments later Henry claims to know that the object he is seeing is a barn. In order to make this claim true, the presuppositions will shift so that this slip of mind counts as ignoring, and therefore Henry's claim does not violate the Rule of Attention.

It is clear that if Lewis (1996) were to define context as he does in his (1979), giving primacy to the Rule of Accommodation, then his (1996) will require nothing more for knowledge than that S believes P, that P is true, and someone claims S knows P. This is unacceptable. In order to preserve Lewis' (1996), one must reject the applicability of the Rule of Accommodation to knowledge ascription, or at the very least make it only one of many factors governing the conversational score. It is unclear how one could do this systematically. Therefore, the prospects for Lewis' contextualism are dim.

Concluding Remarks

DeRose, Stalnaker, and Lewis understand the context of a knowledge attribution as an abstraction on facts about the attributer and putative knower. For them, contexts are transient constructions. Whether the same context ever occurs twice is an irrelevant consideration. Because all three authors fail for reasons associated with their definitions of context, this might lead one to think that a contextualist theory of knowledge needs to rely on a different understanding of context. In this final section, I will sketch what I think a satisfactory definition of context would look like, and list at least two problems it will face.

Recall that Lewis (1996) fails because his rules for knowledge ascription are too elastic, bending to fit the requirements of any knowledge ascription. One way to fix Lewis (1996) is to conceive of contexts as things that exist independently of the facts about the attributer and putative knower. That way, contexts could control how far the rules for knowledge can bend. I suspect that this is not only a potential solution for Lewis (1996), but also the best way to go about constructing a contextualist theory of knowledge.

There are several examples that might tempt one toward this view. For instance, one might think that in a courtroom there are rules for knowledge ascription that exist independently of facts about the particular people that occupy it at any given time. Thus, if a defense lawyer tried to get an eyewitness to admit that she does not know whether the defendant committed the crime because she cannot rule out deception by an omnipotent demon, the witness' knowledge claim is unaffected because in a courtroom such doubts are necessarily irrelevant. In other words, what makes the omnipotent demon hypothesis irrelevant is a fact about court proceedings, not facts about the judge, jury, witness, lawyer, or audience. Similarly, one might think that when deciding if Napoleon knew that the French army would suffer heavy casualties in Russia in a history class, what makes skeptical doubts irrelevant are not facts about the teacher or students, but facts about the nature of history classes.

In this view, each social environment is accompanied by a set of appropriate rules. I will speak of these rules as norms. While it might be possible to change a casual conversation so that skeptical possibilities become relevant, a similar transformation is not possible in a courtroom or laboratory. The mentioning of skeptical possibilities violates the norms of the activity of law or science. In other words, in those contexts the skeptic cannot stretch the pertinent norms so that a skeptical conclusion is inevitable. One advantage of this approach is that, unlike Lewis (1996), it does not make knowledge so elusive that a skeptic only has to whisper into one's ear, "You might be dreaming".

One problem with this approach is that if there are contextual norms governing knowledge ascription then it is hard to determine what they are. This is what I call the access problem. One cannot simply ask participants in a context because few would be able to provide epistemologically satisfactory answers. One cannot merely describe

common behavior because norms are more than descriptions of naturalistic regularities. One needs another method to determine the relevant norms within a context.

Another problem with this view is that it rests on the premise that human activity is rule-governed, which is something questioned by Kripke (1982). Kripke argues that, based on past cases, it is impossible to determine what rule we have followed in the past, and therefore impossible to determine what future action conforms to that rule. This is a direct challenge the concept of normativity. There is a vibrant literature in response to Kripke (1982), but I am uncertain if one can rescue normativity from Kripke's arguments.

I wonder if the reason cases of courtrooms and classrooms work so well is that the judge and teacher are absolute epistemic authorities, possessing the power to decide which doubts are relevant. If we ask whether the judge's and teacher's decisions are themselves bound by rules the answer is not so clear. Judges are regularly overruled by higher courts, the Supreme Court finds it difficult to agree to an interpretation of the Bill of Rights, and teachers differ wildly in the kinds of questions they will entertain. By trying to link a contextualist theory of knowledge to norms similar to those of the courtroom or classroom, there is a chance we are mistaking the firmness of power for the firmness of rules.

We spend most of our lives in situations where there is no clear epistemic authority, so we need another model for norms that govern knowledge ascription. Biology and game theory are examples of disciplines that try to find rules for systems of complex interactions that do not have a clear authority. They might provide the contextualist with a useful model for developing an account of epistemic norms.

In any case, if a contextualist wishes to define contexts as sets of rules that exist apart from facts about its participants, then she must develop a richer understanding of human interactions. I think contextualists can do this only by reevaluating the role of concepts in human practices. This is, to say the least, a much larger burden than contextualists would like to assume.

References

- Brandom, R. (1994). *Making it Explicit*. Cambridge: Harvard University Press.
- Brandom, R. (2000). *Articulating Reasons: An Introduction to Inferentialism*. Cambridge: Harvard University Press.
- Cohen, S. (1988). "How to be a Fallibilist". *Philosophical Perspectives* **2**: 581-605.
- Cohen, S. (1991). "Skepticism, Relevance, and Relativity" in McLaughlin, B. (ed.) *Dretske and His Critics*. Cambridge: Basil Blackwell.
- Cohen, S. (1998). "Contextualist Solutions to Epistemological Problems: Scepticism, Gettier, and the Lottery". *Australian Journal of Philosophy* **76**:289-306. Reprinted in Sosa, Kim (eds.) *Epistemology: An Anthology*. Oxford: Blackwell Publishers Ltd.
- Conee, E.; Feldman, R. (1998). "The Generality Problem for Reliabilism". *Philosophical Studies* **89**:1-29.
- DeRose, K. (1992). "Contextualism and Knowledge Attributions". *Philosophy and Phenomenological Research*. **52**: 913-929.
- DeRose, K. (1995). "Solving the Skeptical Problem". *The Philosophical Review*. **104**:1-52. Reprinted in Sosa, E.; Kim, J. (eds.) (2000) *Epistemology: An Anthology*. Oxford: Blackwell Publishers Ltd. 482-502.
- DeRose, K. (1999). "Contextualism: An Explanation and Defense". In Greco, J.; Sosa, e. (eds.) *Blackwell Guide to Epistemology*. Oxford: Blackwell Publishers Ltd.
- Dretske, F. (1970). "Epistemic Operators". *Journal of Philosophy* **67**: 1007-1023.
- Dretske, F. (1981). "The Pragmatic Dimensions of Knowledge". *Philosophical Studies* **40**: 363- 78.
- Gettier, E. (1963). "Is Justified True Belief Knowledge?". *Analysis*: **23**: 121-123.
- Goldman, A. (1967). "A Causal Theory of Knowing". *The Journal of Philosophy* **64**:357-372.
- Goldman, A. (1976). "Discrimination and Perceptual Knowledge". *The Journal of Philosophy* **73**: 771-791.
- Goodman, N. (1965) *Fact, Fiction and Forecast*. Indianapolis: Bobbs-Merrill.
- Grice, P. (1989). *Studies in the Ways of Words*. Cambridge: Harvard University Press.

- Harman, G. (1973). *Thought*. Princeton: Princeton University Press.
- Hofweber, T. (1999). "Contextualism and the Meaning-Intention Problem". In *Proceedings of the Fifth International Colloquium on Cognitive Science*, Korta, K; Sosa, E; Arrazola, X (eds.) *Philosophical Studies Series*. **79**:93-104.
- Kripke, S. (1982). *Wittgenstein on Rules and Private Language*. Cambridge: Harvard University Press.
- Lewis, D. (1979). "Scorekeeping in a Language Game". *Journal of Philosophical Logic* **8**: 339-59.
- Lewis, D. (1981). "Ordering Semantics and Precise Semantics for Counterfactuals". *Journal of Philosophical Logic* **10**: 217-234. Reprinted in Lewis, D. (1998) *Papers in Philosophical Logic*. 77-96.
- Lewis, D. (1986). *Plurality of Worlds*. Oxford: Oxford University Press.
- Lewis, D. (1996). "Elusive Knowledge". In Sosa, E.; Kim, J. (eds.) (2000) *Epistemology: An Anthology*. Oxford: Blackwell Publishers Ltd. 503-516. Originally in *Australian Journal of Philosophy* **74**:549-567.
- Moore, G. (1962). "Proof of an External World" in *Philosophical Papers*. New York: Collier Books 144-148 Reprinted in Sosa, J.; Kim, K (eds.) *Epistemology: An Anthology*. Oxford: Blackwell Publishers Ltd.
- Nozick, R. (1981). *Philosophical Explanations*. Cambridge: Harvard University Press.
- Pryor, J. (forthcoming). "Highlights of Recent Epistemology". *British Journal for the Philosophy of Science*.
- Schiffer, S. (1996). "Contextualist Solutions to Skepticism". *Proceedings of the Aristotelian Society* **XCVI**: 317-333.
- Stalnaker, R. (1972). "Pragmatics". *Semantics of Natural Language* by Davidson and Harman (eds.)
- Stalnaker, R. (1973). "Presuppositions". *Journal of Philosophical Logic* **2**: 447-57.
- Stalnaker, R. (1984). *Inquiry*. Cambridge, Mass: MIT Press.
- Stanley, J. (2000) "Context and Logical Form", *Linguistics and Philosophy* 23.4: 391-434.
- Stine, G. (1970). "Dretske on Knowing the Logical Consequences". *The Journal of Philosophy* **68**:296-299.
- Stine, G. (1976). "Skepticism, Relevant Alternatives, and Deductive Closure". *Philosophical Studies* **29**: 249-61.

- Sosa, E. (1986). "On Knowledge and Context". *Journal of Philosophy*, **83**:584-585.
- Unger, P. (1975). *Ignorance: A Case of Skepticism*. Oxford: Oxford University Press.
- Unger, P. (1984). *Philosophical Relativity*. Minneapolis: University of Minnesota Press.
- Vogel, J. (1990). "Are There Counterexamples to the Closure Principle?". in Roth, M.; Ross, G. (eds.) *Doubting: Contemporary Perspectives on Skepticism*. Dordrecht: Kluwer.
- Wittgenstein, L. (1969). *On Certainty*. New York: Harper & Row.