

Contextualism and a Puzzle about Seeing

Ram Neta

University of North Carolina, Chapel Hill

Contextualist solutions to skeptical puzzles have recently been subjected to various criticisms. In this paper, I will defend contextualism against an objection prominently pressed by Stanley 2000. According to Stanley, contextualism in epistemology advances an empirically implausible hypothesis about the semantics of knowledge ascriptions in natural language. It is empirically implausible because it attributes to knowledge ascriptions a kind of semantic context-sensitivity that is wholly unlike any well-established type of semantic context-sensitivity in natural language.

More specifically, Stanley 2000 argues that there is no good evidence of any kind of semantic context-sensitivity in natural language over and above the obvious indexicality of personal pronouns like “I” and “they”, token-reflexive expressions like “now” and “here”, and demonstratives like “this” and “that”. All other alleged cases of semantically context-sensitive sentences – cases involving terms such as “flat” or “adequate”, say – are really cases in which the logical form of the sentence in question contains a component that is left unpronounced or hidden. Thus, when we say “the table is flat”, our sentence has the logical form THE TABLE IS FLAT FOR AN F. Contextual variation in the truth-conditions of our sentence is determined solely by variation in the value of F. More generally, variations in the value of this unpronounced component of our sentence show up as variation in the truth-conditions of the sentence, but that’s because they generate variations in which sentence it is. So it’s not that there’s a single sentence that varies in its truth-conditions depending upon the context in which it is uttered. It’s rather that whatever proposition we’re asserting is actually longer than it seems to be, and is different in different contexts. Once the hidden element of the sentence is brought out in the open, we can see that the actual sentence itself is not semantically context-sensitive. Thus, for any complete knowledge ascribing sentence that contains none of the obviously indexical elements mentioned above, that sentence will not be semantically context-sensitive.

Now, why shouldn’t contextualists be happy to say that sentences that are apparently of the form “S knows that p” contain just such unpronounced components in their logical form, and that what the

contextualist describes as the “semantic context-sensitivity” of such sentences is, like the apparent semantic context-sensitivity of ascriptions of flatness, to contextual variation in the value of this unpronounced component?¹ The reason why contextualists should not be happy to say this, according to Stanley, is that sentences apparently of the form “S knows that p” fail a standard test for the presence of unpronounced components of logical form. The test goes like this: embed the predicate of the sentence in a quantifier, and see if the resulting sentence can be read in two different ways corresponding to whether or not the unpronounced component is taken to be in the scope of the quantifier or not. For instance, consider the statement “Everyone went to a local bar”. This can be understood as meaning that each person went to a bar that is local to him or herself (thus, the person in Chicago went to a bar in Chicago, the person in London went to a bar in London, and so on). Or the statement could be understood as meaning that everyone went to the very same bar. This shows that, when we use the modifier “local” in a sentence, we are using an expression the logical form of which is something like LOCAL TO X, where the value of X is within the scope of the universal quantifier on our first reading, and outside the scope of the universal quantifier on the second reading. If we find these two readings, then we have positive evidence of the presence of an unpronounced component. If we don’t find these two readings, then we have no such evidence. And, according to Stanley, we do not find two such readings when we embed predicates of the form “... knows that p” in quantifiers. Thus, there’s no evidence of an unpronounced component in the logical form of sentences of the form “S knows that p”. And so if there’s any semantic context-sensitivity in such sentences, it would have to be due to the presence of an obvious indexical.

I will defend the contextualist against this objection by showing that the kind of unobvious semantic context-sensitivity that the contextualist claims to find in knowledge ascribing sentences is widespread in natural language. Specifically, I will argue that *many* sentences that fail the binding test stated above are nonetheless semantically context-sensitive.

Now, this point has already been ably defended by others, e.g. Ludlow forthcoming and DeRose forthcoming. So my conclusion is not novel. But my argument is novel, and I believe that it is also of some independent interest to philosophers of perception. I will devote most of this paper to constructing a puzzle about seeing. I will then argue that the best solution to this puzzle is a contextualist solution, according to which sentences of the form “S sees x” are semantically context-sensitive. And yet such

sentences, like knowledge ascribing sentences, fail the binding test. So if their semantic context-sensitivity can be explained away as the result of variation in the semantic value of an unpronounced component in their logical form, the binding tests do not reveal this. Finally, I will indicate how an analogous puzzle arises for other relations besides seeing, all of which also fail the binding test. Consequently attributions of these other relations must also be semantically context-sensitive, or else failure of the binding test is not a good way of showing the absence of an unpronounced component in the logical form of a sentence. Either way, it turns out, the recent criticism of epistemological contextualism is wrong: if knowledge ascriptions fail the binding test and yet are semantically context sensitive, that does not distinguish them from many other kinds of sentences in natural language. Of course, Stanley and other critics of contextualism could run my argument in the other direction, and claim that attributions of seeing and other such relations cannot be semantically context-sensitive. But then I hope that my argument has at least considerably increased the cost of their claiming this.

I. A Puzzle About Seeing

I begin by constructing a puzzle about seeing. Suppose that an observer, Alice, is looking at a tomato, and suppose that this takes place under conditions of good lighting and visibility, and that Alice has good vision and is alert and of sound mind. We can represent the situation with figure 1.




Figure 1

Does Alice see the tomato? We would ordinarily say that she does. But of course there are parts of the tomato that she doesn't see. For instance, Alice doesn't see the back of the tomato or the inside of the tomato. In figure 2, I've marked off the portion of the tomato that Alice doesn't see.



Figure 2

Let's call the nearer portion of the tomato "the facing surface". For now, let's not worry about exactly what's included in the facing surface. We'll just say that the facing surface is *whatever is left over* when we mark off whatever portion of the tomato Alice doesn't see.

Now, suppose that we pull the facing surface off of the rest of the tomato, while keeping it perfectly rigid so that it doesn't change shape. And suppose that we move the facing surface closer to Alice, in the direction of her line of sight. That situation is represented in figure 3.



Figure 3

Here, the facing surface occludes the rest of the tomato from Alice's view. In this situation, Alice does not see a tomato. We might of course say that she sees *some* tomato, but then we would be using "tomato" as a mass noun rather than a count noun. And if we said that she still sees "a tomato" in figure 3, we would then be using the phrase "a tomato" to refer to a different, and smaller, object than the object that Alice sees in figure 1. Whatever exactly Alice sees in figure 3, it's not the same thing that we originally took her to see in figure 1. In figure 3, Alice sees only what we've called "the facing surface".

Of course, since it's been severed from the rest of the tomato, it's somewhat misleading to call it "the facing surface", for there is not now any intact object of which it is the surface. But, as long as we keep this fact in mind, we needn't be misled by calling it "the facing surface." We're just using that definite description to pick out a particular object – an object that's part of the tomato in figure 2 but is spatially isolated in figure 3.

Now, suppose that we move the remainder of our erstwhile tomato -- let's call it "the occluded object" -- closer and closer to the facing surface. At some point, the occluded object will touch the facing surface. Before it touches the facing surface, it will still be the case that Alice sees only the facing surface. But once it touches the facing surface, does that make it the case that Alice now sees a numerically distinct object than the one she saw just a moment ago? That's a strange idea. As far as what Alice sees is concerned, *how could it make any difference whether the occluded object happens to be touching the facing surface?* Compare this to the following situation: you look at a concrete wall directly in front of you. Unbeknownst to you, a ladder is leaning against the back of the wall. But do you thereby see the composite object consisting of the wall and the ladder? Whether the ladder touches the wall or not makes no difference to what you see in this case. So too, then, shouldn't it make no difference to what Alice sees whether the occluded object touches the facing surface?

Of course, once the facing surface touches the occluded object, then Alice sees a facing surface with an occluded object behind it. And it may be thought that the two premises "Alice sees a facing surface with an occluded object behind it" and "a facing surface with an occluded object behind it is a tomato" imply the conclusion "Alice sees a tomato", by the extensionality of "sees". But this misunderstands the grammatical role of "with" in the first premise. In the first premise, "with" serves to introduce a relative clause to modify "facing surface". It does not introduce a second direct object for "sees". As it occurs in the first premise then, "Alice sees a facing surface with an occluded object behind it" does not imply "Alice sees the composite object consisting of a facing surface and an occluded object". The sense in which it makes no difference to what Alice sees whether the occluded object touches the facing surface is just this: either way, it is numerically the same object that Alice sees – namely, the facing surface.

Still, we may be reluctant to say that it makes no difference to what Alice sees whether the occluded object touches the facing surface. For if it makes no difference to what Alice sees whether the

occluded object touches the facing surface, then how can Alice see a tomato when she's in the situation represented in figure 1? The only difference between the intact tomato in figure 1 and the object that results when the occluded object touches the facing surface is that, in the latter case, even though the facing surface and the occluded object are touching, they are not attached. But how could it make any difference at all to what Alice sees that the occluded object happens to be attached to the facing surface? Would it make any difference to what you see in the wall and ladder case if the ladder were *glued* to the back of the wall, and not just leaning against it?

You might object that this comparison is unfair: tomatoes, you might say, are *unitary objects*, and so are the kinds of things that we can ordinarily be said to see, whereas composite objects consisting of walls and ladders are not. We might put the point by saying that tomatoes are genuine “wholes” whereas composite objects consisting of walls and ladders are not: thus, we see the former by seeing their facing parts, whereas we do not see the latter by seeing their facing parts. But this reply depends on the assumption that it is only unitary objects that we can truthfully be said to see, and this assumption is false: sometimes we can truthfully be said to see composite objects by seeing their facing parts, and sometimes we cannot be said to see unitary objects, even when we see their facing parts. Suppose that there are two tomatoes directly in front of me, one to the left of the other. In that case, I see the pair of tomatoes. And that is just to say that I see the composite object comprising the pair of tomatoes. If I see the composite object comprising the pair of tomatoes, then why do I not also see the composite object consisting of the wall and ladder? Again, at this moment I see a portion of the Milky Way galaxy, which is a unitary object. The portion that I see is the portion facing me. So why do I not also see the Milky Way galaxy?

Perhaps these responses somehow misunderstand the objector's intended distinction between unitary and composite objects. But if that's so, then the problem with this objection is that it is not sufficiently explanatory: it doesn't tell us what it is *about* tomatoes that makes them “unitary”, i.e., the kinds of things we can ordinarily be said to see. Instead of a tomato, we could just as well have generated a puzzle by thinking about Alice looking at a bowling ball that had just been manufactured by melding two semi-spherical parts together. Also, though it would have made for more complicated diagrams, we could just as well have considered Alice looking at a sofa, or at an automobile, or at Michelangelo's statue *David*. What do tomatoes, bowling balls, sofas, automobiles, pairs of tomatoes, and the statue *David* all have that

the Milky Way galaxy, and composite objects consisting of walls and ladders, do not have? Let's grant that the former are all things that we can *ordinarily* be said to see, whereas composite objects consisting of walls and ladders are not, and neither is the Milky Way. But this is simply to point out a fact about how we typically regard certain objects. It doesn't explain why we so regard them, or what would make us right in so regarding them.

What's wrong with our treating the tomato as a composite object? An Aristotelian realist about substances would claim that tomatoes *simply are* unitary objects.² So are bowling balls, sofas, automobiles, and Michelangelo's statue *David*. But what about pairs of tomatoes? And why not the Milky Way galaxy? Even if it could answer these questions, such Aristotelian realism would have the highly implausible consequence that what Alice sees at a given moment depends constitutively upon contingent facts about the history of the things in front of her. What does Aristotelian realism say, for instance, about a wall with a ladder leaning against it? Well, if the wall and ladder jointly constitute an artwork, then the result is a unitary object, and so it's something that Alice sees. But what if it's not an artwork? Then maybe it's composite, and so not something that Alice sees. But how can it matter to what Alice sees whether the objects in front of her were arranged by human artifice or not? If Alice sees a particular thing, then – I would think – she sees that thing no matter whether the object is the product of human artifice.

Mereology won't help us to distinguish tomatoes, bowling balls, and sofas on the one hand from wall-ladders on the other hand. The orthodox view of mereological composition is the view of *unrestricted composition*, on which any two objects have a fusion. And while restricted principles of composition have been proposed, none of those principles (so far as I know) distinguishes sofas and the statue of *David*, on the one hand, from wall-ladders, on the other.

Again, it will not help us to point out that facing surfaces of tomatoes are reliably correlated with whole tomatoes, whereas walls aren't reliably correlated with wall-ladders. Even if this is true, it still won't solve our puzzle. Suppose that Martians start replacing tomatoes throughout the world with tomato surfaces, and they eventually succeed in destroying the reliable correlation between tomato surfaces and tomatoes. Still, if Alice happens to be one of the lucky few who's looking at a whole tomato, don't we want to say that *she sees the tomato* (whether or not she knows that that's what she's seeing)? In this case, the general failure of reliable correlation between facing surface and tomato makes no difference to what

Alice sees. Even without there being a reliable correlation of facing surfaces and tomatoes, Alice still sees *the tomato*. This comes out clearly when we think about Carl Ginet's case of barn façade county, made famous by Goldman 1976. When you're looking at the only real barn in barn façade county, then, even though you don't know that there's a barn in front of you, nonetheless you still see a barn, and not just a barn façade.

Notice also that, in *some* possible circumstances, it is natural to say that we see objects consisting of walls and ladders. Suppose that we come upon an artwork entitled "Wall and Ladder": it consists of a wall with a ladder attached to the back of it. (It is the fusion of the two. If the artist's intentions are relevant here, we may suppose that an ontologically-minded artist intended the artwork to be the fusion of the wall and the ladder.) When you look at the artwork, you can naturally be said to see the artwork consisting of the wall with a ladder attached to the back of it. So why isn't it equally natural to say that you see the composite object consisting of the wall and ladder in other circumstances? And why do we have to think up special circumstances under which someone can naturally be said to see the wall-ladder, but we needn't think of special circumstances under which someone can naturally be said to see the tomato?

And so, regarding the issue of what Alice sees, it is not easy to understand exactly how it could make any difference whether the occluded object happens to be touching, or attached to, the facing surface. It may then seem as if, even when there is an intact tomato in front of Alice, she still doesn't see that tomato. She sees only its facing surface. But, the example I've chosen is representative: what holds true of Alice in this case holds true of all of us quite generally. *Thus, no one ever sees the ordinary physical objects that we ordinarily take ourselves to see.* We never see tomatoes, books, people, buildings, and so on. We see only their facing surfaces.

This conclusion deals a heavy blow to common sense. But we might be tempted to soften the blow by means of the following line of thought: our ordinary ascriptions of seeing are really just loose talk. What we *really* see at any given moment – as students of vision should know – are simply those bunches of molecules that have some causal impact on our retinas at that moment. These bunches of molecules are just the "facing surfaces" of which I've been speaking. And so it should be no surprise if those are all that we ever really see.

Now, this view faces a dilemma when it confronts the question: do facing surfaces have any thickness at all? If the answer is “yes”, then the puzzle about seeing arises all over again for facing surfaces, and it turns out that we never really see them either, but we see only *their* facing surfaces. If the answer is “no”, then we avoid the puzzle about seeing at the cost of claiming that our retinas are causally affected by distal objects with no thickness. But a distal object with no thickness cannot contain any molecules, since molecules themselves have thickness. So how can a distal object that contains no molecules (assuming there are such things) have any causal impact on our retinas? There is nothing commonsensical about this idea, and so the suggestion above fails to soften the blow to common sense. In short, common sense has no way to accommodate the conclusion that we never see any of the ordinary physical objects that we ordinarily take ourselves to see.

How then can we plausibly avoid this conclusion? I will call this question – the question of how we can plausibly avoid the repugnant conclusion just stated – our “puzzle about seeing”. Recall that what led us to this puzzle was the intuitively compelling thought that *it can make no difference to what Alice sees that there happens to be an occluded object behind, touching, or attached to the facing surface that she sees*. I’ll call this the “No Difference Principle”.

If we could find a way to resist the intuitive force of the No Difference Principle, then we could solve our puzzle about seeing.

We can sum up our puzzle about seeing in the following argument:

- (1) When Alice is in figure 2, she sees the same facing surface that she sees when she’s in figure 3 (even if she sees more than merely that facing surface when in figure 2).³
- (2) There is no difference between what Alice sees when she’s in figure 1 and what she sees when she’s in figure 2.
- (3) Therefore (from 1, 2), when Alice is in figure 1, she sees the same facing surface that she sees when she’s in figure 3 (even if she sees more than that when in figure 1).

(4) If Alice sees a facing surface when she's in figure 1, then she *also* sees a tomato when she's in figure 1 only if the presence of the occluded part of a tomato behind (touching, attached to) the facing surface makes a difference to what she sees.

(5) Therefore (from 3, 4), Alice sees a tomato when she's in figure 1 only if the presence of the occluded part of a tomato behind (touching, attached to) the facing surface that she sees makes a difference to what she sees.

(No Difference Principle) If there happens to be an occluded object behind (touching, attached to) the facing surface that Alice sees, that makes no difference to what Alice sees.

(Conclusion) Alice does not see a tomato in figure 1. (from 5, No Difference Principle)

So the No Difference Principle, in conjunction with obvious facts about the situations depicted in figures 1, 2, and 3, generates a conclusion that we must avoid. Our puzzle about seeing is this: how do we avoid that conclusion?

The argument above is valid, so it seems that the only way to avoid the conclusion is by denying the No Difference Principle. But then how to explain away the plausibility of that principle? We might try to explain away its plausibility by claiming that it trades on two senses of the phrase "what Alice sees": in one sense, the phrase refers to Alice's visual experience, which remains constant across the situations described, whereas in the other sense, the phrase refers to the external object that Alice sees, which does not remain constant across those situations. But this attempt to explain away the plausibility of the No Difference Principle fails, for the plausibility of that principle is preserved even when we express it in the following more cumbersome but less ambiguous sentence: If there happens to be an occluded object behind (touching, attached to) the facing surface that Alice sees, that makes no difference to the issue of which external thing(s) Alice sees. This principle, so clarified, is still very hard to resist. That's precisely what the example of the wall and ladder was intended to show.

II. A Contextualist Solution to the Puzzle about Seeing

So our puzzle about seeing seems to be this: we have a seemingly sound argument for an obviously false conclusion. What has gone wrong? A contextualist solution to this puzzle will claim that statements of the form “S sees x” are semantically context-sensitive, and that going through the steps of the aforestated argument involves some tacit shift in the semantically relevant features of the context. But where and how might the alleged context-shift take place? Our answer to that question will depend upon the details of our contextualist theory of “S sees x”.⁴ For present purposes, I need only state a portion of my own view of the matter. I turn to that now.

Seeing a thing requires at least seeing some part of that thing. I cannot see a thing without seeing any part of it at all. But does seeing any part of a thing *suffice* for seeing that thing? No. Right now, I see a part of the universe. But I don’t see the universe. Again, when I look down at the ground from my present height of about six feet, I see a part of the planet Earth, but I don’t see the planet Earth. And again, when I’m sitting in my office, I see an interior part of the building that my office is in. But I do not see the building. So, while seeing a part of x is necessary for seeing x, it is not sufficient.

What else – besides merely seeing part of x – is required for seeing x? Let’s consider some examples: Typically, it suffices for my seeing you that I see your face, but it does not suffice for my seeing you that I see a single strand or your hair sticking out from under the blanket. Typically, it suffices for my seeing the field that I see a sufficient portion of its grassy surface, but it does not suffice for my seeing the field that I see a small piece of the dirt lying between blades of its grass. Typically, it suffices my seeing the sofa that I see a sufficient portion of its outer surface, but it does not suffice for my seeing the sofa that I see one of the springs inside it (through some sort of video device installed inside the sofa). In general, seeing x requires not merely seeing a part of x, but seeing a sufficient and characteristic part of x. But how much is sufficient, and which parts are characteristic?

The answers to these questions depend in part, I claim, upon the context in which the seeing ascription is made. To illustrate: Suppose that Jones works for a ratings board that rates stage plays for sexual content. As he’s watching a performance of the play, he is considering the question whether a viewer of the play can see certain body parts of the actors on stage. Another member of the audience – Smith – does not work for a ratings board. Smith is the director of the play, and is watching the

performance and wondering what advice he can give the actors for their next performance. Smith has a particular interpretation of the script, and believes that it's crucial to the aesthetic success of the play that the audience be able to see certain body parts of the actors on stage. Now we might pick a particular body part of a particular actor and ask: at a given moment in the performance, can the audience see that body part? Jones might correctly answer this question in the affirmative when reporting back to the ratings board, even though Smith correctly answers this question in the negative when debriefing the actors after the performance. Jones might correctly say to the ratings board "We have to give this play an X rating, since you could see the actor's private parts", while Smith might correctly say to the actor after the performance "You were turning away from the audience so that they couldn't see your private parts." Take out of context, their seeing ascriptions are apparently incompatible, but it seems that, given the different contexts in which they are made, they could both be true. If we claim that one or the other seeing ascription was false, then we will be in the difficult position of having to say which one was false, and why. I suggest that, instead, we reject the appearance of incompatibility, and allow that the two apparently incompatible seeing ascriptions are not actually incompatible.

Inspection of more or less ordinary cases provides some support for this contextualism about "S sees x". But it seems to me that some powerful additional support for this contextualist view is provided by its ability to solve our puzzle about seeing. I'll sketch the solution now: Recall that, when we drew figure 2, we left it open exactly what part of the tomato should be included in "the occluded part" and what part of it should be included in "the facing surface". But it turns out that this is crucial. Whether or not Alice sees a tomato depends upon whether or not what we mark out as the "facing surface" constitutes a sufficient and characteristic portion of the tomato. If the facing surface does not, in our context, constitute a sufficient and characteristic portion of the tomato, then we cannot truthfully credit Alice with seeing the tomato. What we can truthfully say about what Alice sees depends not just on the physical facts about Alice and her environment, but also depends upon whether we, the *ascribers* of the seeing, treat the facing surface marked out in figure 2 as a sufficient and characteristic portion of the tomato. If it does constitute such a portion, then we can truthfully say that Alice sees the tomato. Otherwise, we cannot.

In a context in which we select a boundary between facing surface and occluded part that renders the facing surface insufficiently large for Alice to count as seeing the tomato, the argument given above

(from 1-5 and the No Difference Principle) is sound, and the conclusion that Alice doesn't see the tomato is true. But in most ordinary contexts, we don't select any boundary between its facing surface and its occluded part. In those ordinary contexts, premise 1 of the argument above isn't true, for it employs the phrase "facing surface", which suffers from reference failure. And so in such contexts, the argument isn't sound.

The example of Alice and the tomato is representative, and so I conclude that sentences of the form "S sees x" are semantically context-sensitive. What object we can truthfully say that you see depends not just on the physical facts about you and your environment, it depends also on what *we* – the *ascribers* of the seeing – treat as a sufficient or characteristic part of an object for the purposes of our seeing ascription (or the more inclusive stretch of thought or speech of which that seeing ascription is a part).

III. Conclusion: Generalizing this Result

The foregoing argument for the semantic context-sensitivity of sentences of the form "S sees x" applies equally well to various other relational sentences, including sentences of the form "S feels x", "S tastes x", and others. I leave it to the reader to construct, by analogy with our puzzle about seeing, puzzles about feeling and tasting. Each of the analogous puzzles can be solved only by means of a perfectly analogous contextualist solution. These various other relational ascriptions are all semantically context-sensitive.

Now we might try to avoid this conclusion by claiming that the relational ascriptions that we've considered are all incomplete. For instance, we might say, sentences that we typically think of as having the form "S sees x" really have a component that's not indicated in that schema, and the same goes for the other relational ascriptions that we've just canvassed (e.g. "S feels x", "S tastes x"). But the binding tests that I described above provide no confirmation (at least none that I can find) for the hypothesis that such sentences contain an unpronounced component in their logical form. So this way of avoiding contextualism about seeing commits us to claiming that the failure of sentences of the form "S knows that p" with respect to the binding tests does not tell against the presence of an unpronounced component in the logical form of such sentences. And this is enough to defend epistemological contextualism against the

argument of Stanley 2000: the contextualist can claim that Stanley's argument fails to tell against the incompleteness of "S knows that p".

It turns out then, that the kind of hidden semantic feature that contextualists have claimed to find in sentences of the form "S knows that p" are not linguistically unprecedented at all. Semantic context-sensitivity is a widespread feature of sentences in natural language, including sentences that contain no obvious indexicals or ellipses. There should be nothing surprising if it is a feature of knowledge ascriptions as well. If this is right, then Stanley 2000 has given us no good reason to forego the advantages of contextualist solutions to skeptical puzzles.⁵

Works Cited

- Black, Max. 1964. *Philosophy in America*. Ithaca, NY: Cornell University Press.
- Clarke, Thompson. 1964. "Seeing Surfaces and Physical Objects" in Black 1964.
- Cohen, Stewart. 1999. "Contextualism, Skepticism, and Reasons" in Tomberlin 1999.
- DeRose, Keith. Forthcoming. "The Ordinary Language Basis for Contextualism and the New Invariantism."
- Goldman, Alvin. 1976. "Discrimination and Perceptual Knowledge." *Journal of Philosophy*
- Ludlow, Peter. Forthcoming. "Contextualism and the New Linguistic Turn in Epistemology" in Preyer and Peter forthcoming.
- Preyer, Gerhard and Peter, Georg, eds. Forthcoming. *Contextualism in Philosophy: On Epistemology, Language, and Truth*. Oxford: Oxford University Press.
- Schaffer, Jonathan. 2004. "From Contextualism to Contrastivism." *Philosophical Studies* 119: 73 – 103.
- Stanley, Jason. 2000. "Context and Logical Form." *Linguistics and Philosophy* 23: 391 – 434.
- Tomberlin, James, ed. 1999. *Philosophical Perspectives 13: Epistemology*. Malden, MA: Blackwell Publishers.

¹ Such a view is adopted by Schaffer 2004, and entertained by Cohen 1999.

² Thanks to Brian Weatherson for bringing this point to my attention.

³ Someone might object: "The thing that Alice sees in figure 3 – the thing that we've called 'the facing surface' – is not the surface of a tomato, since there is no intact tomato for it to be the surface of. But the thing that she sees in figure 1 is the surface of a tomato. And so the thing she sees in figure 3 cannot be identical to anything that she sees in figure 1." While I believe that this objection rests on a false presupposition concerning the identity conditions for things that persist through time, I will not attempt to argue for that here. Rather, I will just point out that the argument of the text can be reformulated equally well in terms of indiscernibility rather than identity.

⁴ My contextualism about "S sees x" should not be confused with the view defended by Clarke 1964, according to which the correct answer to a question of the form "How much of x do you see?" depends upon the context in which the question arises.

⁵ I am grateful to Dorit Bar-On, James Cargile, Jonathan Cohen, Wayne Davis, Rick Gallimore, Mark Greenberg, Elizabeth Harman, Avram Hiller, William Lycan, Eric Marcus, Jonathan Schaffer, Cara Spencer, Jonathan Vogel, and Brian Weatherson for helpful comments and questions concerning earlier versions of this material.