

Inexplicit Thoughts

Christopher Gauker

Abstract: It is often assumed that, though we may speak in sentences that express propositions only inexplicitly, our thoughts must express their propositional contents explicitly. This paper argues that, on the contrary, thoughts too may be inexplicit. Inexplicit thoughts may effectively drive behavior inasmuch as they rest on a foundation of imagistic cognition. The paper also sketches an approach to semantic theory that accommodates inexplicitness in mental representations as well as in spoken sentences.

Keywords: Communication, conceptual thought, imagistic cognition, context-relativity

Index items: accessibility criteria, anaphora, communication, computational theory of mind, conceptual thought, context, context-relativity, domain of discourse, ellipsis, explicitism, extended mind hypothesis, imagistic cognition, logical validity, pragmatics, rational choice, semantics, situation, subsentential speech

Biography: Christopher Gauker is Professor of Philosophy at the University of Cincinnati. He is the author of *Thinking Out Loud: An Essay on the Relation between Thought and Language* (Princeton 1994), *Words without Meaning* (MIT 2003), *Conditionals in Context* (MIT 2005) and *Words and Images: An Essay on the Origin of Ideas* (Oxford 2011).

Inexplicit Thoughts

Christopher Gauker

1. Explicitism

It is commonly assumed that the vehicles of occurrent, conceptual thought must be, in some sense, fully explicit. The propositional content of an utterance, everyone agrees, depends not only on the meaning of the sentence uttered but also on various features of the situation in which it is uttered. But the propositional content of a particular conceptual thought, it is supposed, cannot likewise depend on the situation in which it occurs. Or it can depend on the situation in only a few of the ways in which the content of an utterance does. Call this view *explicitism*. I will argue that explicitism is false. The content of a conceptual thought can depend on the situation in which it occurs in all of the ways that the content of an utterance of a natural language sentence can do.

Explicitists generally deny that we think in natural languages. But because I cannot write sentences of Mentalese on paper, I will here and throughout use English sentences as if they were sentences of Mentalese. In spoken English we might say something like “Everyone likes that”. In the situation in which this sentence is uttered, the proposition the utterance expresses might be: *Everyone now at this party enjoys the taste of that*, where *that* equals the tapenade that the speaker is pointing at. But according to explicitism, the thought that bears that proposition will never be merely a token of the Mentalese sentence, “Everyone likes that”. The thought will have to be a token of the Mentalese sentence, “Everyone now at this party enjoys the taste of that”. For some reason, explicitists are willing to allow that thoughts may contain demonstratives like

“now” and “that”. But the set of people now at this party cannot be expressed in Mentalese with the unadorned quantifier “everyone”.¹

Let us say that a representation type expresses a proposition *explicitly* if and only if that proposition can be read off the representation type in light of the semantic properties of the language of representation together with an assignment of referents to demonstrative expressions; to identify that proposition as the proposition expressed, it is not necessary to consider features of the situation in which the token occurs not represented in the representation by virtue of its type, except insofar as doing so is necessary in order to assign a reference to demonstrative expressions. (Here I speak of *expression* as a relation between a representation and its content.) Setting aside difficult questions about “representation type” and “semantic properties”, we may define explicitism as follows: The explicitist holds that a thought is always a token of representation type that *explicitly* expresses the proposition that is the content of the thought. In this sense, thoughts are always explicit.

Elsewhere, I have argued against explicitism by arguing that sometimes there is no plausible account of the explicit thought underlying a speaker’s words. I will review that style of argument below. But my main objective in this paper is to argue in a different way, by showing, through examples, that explicitness in thought is not required. Thinkers can solve their problems by means of inexplicit thoughts, and interlocutors can track their conversations by means of inexplicit thoughts.

If explicitism were true, it would provide a strong argument against the identification of occurrent conceptual thoughts with tokens of sentences in natural language. One could try to combine explicitism with the thesis that spoken language is

the very medium of conceptual thought by arguing that we think in sentences of natural language that are much more explicit than the sentences we normally speak. But, for various reasons, that will be a hard argument to make. So to maintain that spoken language is the very medium of conceptual thought, we are pretty much forced to abandon explicitism. Thus, in arguing against explicitism, I remove one barrier to identifying occurrent conceptual thoughts with acts of speech in the very languages we speak.

Likewise, explicitism undergirds a common model of communication, according to which linguistic communication is a process by which a speaker encodes his or her thought into words and the hearer decodes the speaker's words back into thought (this "thought" becoming a "belief" only if the hearer accepts what the speaker says). If explicitism is true, then, since the thoughts that words express are always explicit, while the words that express them may be inexplicit, it will be very tempting to suppose that a hearer can understand a speaker's words only by writing out in his or her own thought the explicitly represented content of the speaker's underlying thought. Thus, in arguing that explicitism is false, I remove an important motivation for this model of communication.

Explicitism does not say that the content of a thought about a thing includes all that the thinker happens to know about that thing. The explicitist expects only that the propositional content of each thought, limited though it may be, will be explicitly expressed by its vehicle. If someone thinks, "The pilot of this airplane is experienced", then, even though the thinker happens to know that the pilot is British, the thought may be perfectly explicit, because the proposition that the pilot is British may not be part of the propositional content of this particular thought. Moreover, explicitism, as I define it,

is exclusively a thesis about *occurrent* conceptual thoughts. Explicitism could be extended to cover what we might call *stored* beliefs, but I will not try to extend it in that way. Explicitists need not deny that we have beliefs the content of which is only implicit in our occurrent thoughts and stored beliefs.

Explicitism is a thesis specifically about *conceptual thought*. Here I define *conceptual thoughts* as thoughts of a kind that are either identical to or *expressed* by tokens of natural language sentences. (Here I speak of *expression* as a relation between an external representation and an internal representation.) For my purposes, it is important to leave open the possibility that there are other, nonconceptual varieties of thought, such as imagistic thought, and to allow that these other varieties may operate independently from conceptual thought. I will grant that conceptual thoughts may have to be more explicit in some ways than spoken sentences. For example, thoughts may wear the resolution of anaphoric ambiguities on their sleeves. But I will draw a distinction between these kinds of explicitness and the kind of explicitness required by explicitists.

2. Two arguments for explicitism

A first argument for explicitism goes like this: The meaning of an utterance cannot always be determined from the semantic meaning of the sentence uttered. Pragmatics has to come into play as well. But giving pragmatics a role just means identifying what the speaker intended or had in mind in speaking. If we are to avoid a regress, we cannot suppose that the meaning of the speaker's thought likewise depends on the content of some other underlying thought in the speaker. So the meaning of a thought cannot likewise depend on pragmatics. Its meaning must be entirely determined by the semantics

of the thought type of which the particular thought in question is a token (and an assignment of referents to demonstratives).²

To clarify this argument, we should distinguish between “pragmatics” considered as a process that the hearer must engage in in order to understand a speaker’s utterance and “pragmatics” considered as a discipline that aims at explaining in a general way the conditions under which a given utterance expresses a given proposition. The argument presupposes that the process by which a hearer understands a speaker’s utterance – pragmatics in the first sense – consists in identifying the speaker’s underlying thought. But proponents of this argument may draw encouragement from a conception of pragmatics as a discipline, according to which a pragmatic theory explains in a general way how the proposition that an utterance expresses depends on the content of the speaker’s underlying thoughts.

A second argument for explicitism goes like this: The cognitive work a thought can do depends on the structure of its concrete vehicle in the brain. If there are aspects of the content of a thought that are due not to the concrete structure of the vehicle of thought, but due only to the context relative to which it is interpreted or evaluated, then those aspects of content cannot correspond to any efficacious feature of the thought at all, and there is no reason to speak of them as far as our explanations of behavior are concerned. But when we consider the aspects of the content of a spoken sentence that might be supposed to be due to context, we find that those aspects of the content of thoughts definitely are important for our explanations of behavior. The host makes more tapenade because everyone *at the party* likes it, not because everyone *in the universe*

likes it. So those aspects of the content of thought must be reflected in structures of the vehicle that bears them.³

3. Some concessions

I do not wish to deny that the sentences we think in may often be more explicit than the sentences we speak. Of course, we may sometimes deliberately edit down what we think before speaking, and we do it for various reasons, ranging from economy to politeness. Moreover, there may be certain questions of meaning that have to be settled in thought, though the form of words that may be used to express the thought does not on the surface indicate how those questions are to be settled.

Consider first the case of anaphoric cross reference. In thought, though not in speech, there may have to be some explicit indication of which pronoun picks up the reference of a prior noun. Compare:

- (1) a. John₁ took the pizza from the pizza delivery man₂. He₁ paid him₂.
- b. The pizza delivery man₁ handed John₂ the pizza. He₂ paid him₁.

Of course, there is no reason to think that there are numerical subscripts in the brain. Nonetheless, there might be in the Mentalese equivalents of these sentences something that indicates whether the first “he” refers to the same thing as “John” or to the same thing as “the pizza delivery man”. Against this, it might be suggested that thoughts are structured in such a way that we need only one designator per referent. But that is unlikely, because sometimes we discover identities; so there has to be some way the mind can mark in thought the fact that two designators co-refer, for instance, by thinking explicit identities of the form $a = b$.

Second, consider the case of syntactic ambiguity, as in the following classic example:

(2) Visiting relatives can be boring.

We can parse that in either of two ways, roughly:

[[t_1 [visiting relatives]]₂ [can [t_2 be boring (to) t_1]]]

[[[t_1 visiting t_2] [relatives₁]] [can [t_1 be boring (to) t_2]]]

Here we have not only alternative ways of resolving co-reference but also alternative ways of parsing the phrase “visiting relatives”. When a speaker speaks such a sentence, he or she presumably has one or the other of these two parsings in mind. In other words, the thought expressed must have one or the other of these two parsings built into its intrinsic structure and in that respect will be more explicit than the sentence spoken.

Third, there is the case of ellipsis. Ellipsis is the fact that a sentence may omit certain words, otherwise grammatically required, due to the presence of correlated material elsewhere in the sentence. For example, in the following sentence, the phrase “he flattered” is elided at the end of the sentence.

(3) He flattered someone, but I don’t know whom.

In languages with a greater variety of case markings than English, we can find evidence that in some sense the elided material is present in the mind’s representation of the sentence, because it affects the case of the word corresponding to “whom” (Merchant, forthcoming).

From the hearer’s point of view, utterances of any of these sorts of sentences require some interpretation on the basis of the situation in which the utterances take place. In the case of “He paid him”, the hearer has to figure who was most likely to be

paid, John or the pizza delivery man. In the case of “Visiting relatives can be boring”, the hearer has to figure out whether the subject of “can be boring” is the relatives or the act of visiting. In the case of ellipsis, the hearer may have to make some inferences about what material has been elided. But in none of these cases, does the process of representing the content of the sentence explicitly involve finding additional lexical material not already contained somewhere in the sentence on the basis of features of the situation in which the utterance occurs.

The explicitist expects that thought will be explicit in ways that go beyond these. A thought cannot lack open class expressions, not found elsewhere in the thought, that would be found in any explicit expression of its content. For example, “Everyone likes that” cannot stand in for “Everyone *at this party* likes that”, “Dumbo is small” cannot stand in for “Dumbo is small *for an elephant*”, and “Tipper is ready” cannot stand in for “Tipper is ready *to go to the convention center*”. My claim is that explicitism goes too far in supposing that thought is always explicit in these ways. Let us say that a sentence *S* (whether of spoken language or Mentalese) is *lexically inexplicit* if and only if it is not explicit in the sense I defined at the start *and* for each situation in which an utterance (or tokening) of *S* would express a proposition *p*, any sentence that explicitly expressed *p* would contain lexical material not in *S*. My claim is that thoughts can be lexically inexplicit in all of the ways that sentences of spoken language can be.

In thus limiting my critique of explicitism I do not wish to maintain that thought must always be explicit with respect to the resolution of anaphora, syntactic structure and ellipsis. Thoughts can be confused in all sorts of ways, including not being perfectly resolved with respect to anaphora, syntactic structure and ellipsis. My disagreement with

explicitism pertains to thoughts that are at least not confused. Even perfectly clear thoughts may be lexically inexplicit.

4. The *prima facie* case against explicitism

Many years ago I published a paper (1997) in which I attempted to show that thoughts are not explicit with respect to domains of discourse. In that paper, I told a story about a Peruvian goatherd who witnesses a bright falling star one evening during an annual celebration in his village. The next day he is out in the hills tending his goats and has occasion to say (in his own native language), “Everyone saw the falling star”. The goatherd’s village is so remote that no one living there has ever encountered an outsider. This particular goatherd is so short on curiosity that it never occurs to him to wonder whether there might be other people elsewhere in the world. So when he says, “Everyone saw the falling star”, it is not credible that he modifies the quantifier in thought by adding “who lives in the village” to form the quantifier “Everyone who lives in the village”. Nonetheless, the content of his thought is clearly that *everyone who lives in the village saw the falling star*.

In my paper, I considered a reply, according to which what I have translated as “everyone” is an expression of the goatherd’s native language that literally means “everyone who lives in the village”. On that interpretation, even the words he speaks explicitly express the propositional content of his thought. But this interpretation is not plausible. To see that it is not, suppose that on that same day a philosophical friend asks him, “Do you think there might be other people beyond those distant mountains?” For the first time, our goatherd considers the question and concludes that, yes, probably there are

other people somewhere, a conclusion that he expresses by saying, “Not everyone lives in the village”. If “everyone” in his language just means “everyone who lives in the village”, then his answer would mean, “Not everyone who lives in the village lives in the village”. But it would be unreasonably uncharitable to interpret the goatherd as expressing such a claim.

In reply, it might be said that *for the goatherd*, before he takes up the philosopher’s question, the proposition that *everyone saw the falling star* just is the proposition that *everyone who lives in the village saw the falling star*. So from the goatherd’s point of view, it might be said, the words “Everyone saw the falling star” are a perfectly explicit expression of his thought.⁴ On the contrary, it is no defense of explicitism to say that the goatherd’s thought counts as explicit so long as we pretend that we are unaware of the existence of people outside the goatherd’s village. Speaking our own language, not ignoring what we know, the proposition that the goatherd expresses is precisely this: Everyone who lives in the village saw the falling star. That is not a proposition that his sentence or his thought explicitly expresses.

During the time when the goatherd does not add “who lives in the village” to his thought, he has never contemplated the possibility that there might be other people. Perhaps the explicitist could allow, without sacrificing much, that in such ways thought is inexplicit, because what is not explicitly expressed is only a feature of the content of thought that represents an invariable feature of the environment.⁵ We can grant that the goatherd thinks of *everyone in the village* by means of the quantified noun phrase “everyone” just because the population of the village is a constant. The trouble for this answer is that, in fact, the population of the village is not a constant. People are born and

people die. The reason it was implausible that the goatherd's thought must contain a qualifier, attached to the quantified noun, such as "who lives in the village", was not just that he gave no thought to existence of people elsewhere in the world but that the variability of the population was something he had no reason to think about in the situation in which he spoke.

Similarly, we can take any of the kinds of context-relativity we find in spoken sentences and tell a story in which a thought may be inexplicit, just by imaging a situation in which the thinker has no special reason to add in thought the qualifications that make his or her thought fully explicit. Consider the following scenario from the modern American economy. Buddy is a chauffeur for Mrs. Winkelvoss. His only job is to take Mrs. Winkelvoss shopping. She wants to go shopping every day and is always ready to go sometime between 10 and 11 a.m. Parking nearby the entrance to her building, he waits for a telephone call from the doorman. When the doorman calls, the doorman says, "Mrs. Winkelvoss is ready". What Mrs. Winkelvoss is ready for is to go shopping. Buddy hears and understands what the doorman says. But he does not have to think, "Mrs. Winkelvoss is ready to go shopping". He can think simply, "Mrs. Winkelvoss is ready". He does not have to add "to go shopping" in his thought, because what she is ready for under these circumstances, at this time of day, is always the same thing: to go shopping. And yet, the propositional content of his thought on each occasion is that Mrs. Winkelvoss is ready *to go shopping*. Of course, Buddy knows in some sense that Mrs. Winkelvoss is going shopping. But going shopping need not be explicitly represented by the representational components that make up Buddy's thought.

5. Conversing without explicit thought

I now wish to answer the first argument for explicitism by arguing that it is quite possible for two people to carry on an intelligible, productive conversation by means of inexplicit sentences, though the underlying thoughts are inexplicit as well. In discussing this issue, I will play along with the idea that in a conversation each utterance is underlain in the utterer by a “thought”, and I will speak as if this underlying thought is an event distinct from the utterance that expresses it. In reality, I do not assume that occurrent conceptual thoughts need be distinct from overt utterances. On the contrary, in refuting explicitism, I remove one barrier to identifying occurrent conceptual thoughts with acts of speech in the very languages we speak (either overt or covert).

As others before me have emphasized, the structure of the environment might in various ways support the use of inexplicit thoughts in cognition.⁶ One way in which the environment can obviate the need for explicitness in thought is by, so to speak, instantiating routines. If a man who catches lobsters for a living comes back home from a day of checking his traps and reports to his wife, “I caught a lot”, she does not, in order to understand him, have to think to herself, “He caught more *than 30 lobsters today*”. She understands that he is talking about catching *lobsters* because lobsters is what he catches every day. She understands that he is talking about *that day’s* catch because what her husband reports when he comes home is always the catch of that same day. She understands that he caught more than 30 lobsters, because he never reports catching “a lot” unless he has caught more than 30. The example of Mrs. Winkelvoss’s chauffeur also illustrates the fact that routine can obviate the need for explicitness.

Another way in which the environment can obviate the need for explicitness is by revealing the goals of activity and thereby imposing a criterion of relevance. Suppose we are filling buckets with oats to feed to horses, one for each horse in the stable. As I am pouring oats from a bag into the first bucket, you tell me, “That’s enough”. At that point, I stop pouring oats into that bucket. But I still take another bucket and pour oats into it. I understand that you mean that I have poured enough oats into this bucket and not that I have filled enough buckets. Because I understand our goals, I understand that I should stop pouring oats into this bucket but should go on pouring oats into more buckets. I do not have to think, “The amount of oats I have poured into this bucket is as much as I should pour”.

Another way, perhaps, in which the environment can obviate the need for explicitness is by providing standards for gradable adjectives. Suppose someone says, “These tomatoes are really sweet”. I may understand her even though I do not, in advance, have any conception of what counts as sweetness in a tomato and would not be able to distinguish between sweet and nonsweet tomatoes (because, by comparison with paradigmatically sweet things, such as sugar candy, no tomatoes are sweet). I understand her, provided I am willing to grant her authority, by understanding that a “really sweet” tomato is a tomato that has the flavor of the tomatoes she is referring to. Or if someone tells me, “A tall man will meet you at the airport”, I will understand him perfectly well although I might not know until I get there what counts as tallness in the land I am traveling to.

One might have expected a different sort of appeal to the environment than the ones I am making. One might have thought that thoughts can be inexplicit because the

structure that is missing from an inexplicit thought may be found in the environment instead. For instance, one might have supposed that a thought that reads simply “Everyone likes *that*” might bear the proposition that *everyone at the party likes that* because the part about being at the party is borne by the environment instead of being borne by the vehicle of thought. But to say that one would have to make sense of the idea that the environment can bear conceptual content (conceived as components of propositional content). But I do not know how to make sense of that, and so I have not proposed to appeal to the environment in any such way. More generally, I have not committed myself to any version of what is known as the “extended mind hypothesis” (Clark and Chalmers 1998).

Against my purported examples of inexplicit thought, it may be objected that there must be something going on in the thinker’s head that constitutes the understanding I have been talking about. So if the lobster man’s wife *understands* that her husband is speaking of lobsters, there has to be something going on her head that, in her environment, constitutes her thinking of lobsters in interpreting her husband’s words. This is not a denial of externalism, the thesis that the content of a thought depends on the character of the environment in which the thinker is embedded. We can agree that whether “water” refers to H₂O or XYZ depends not only on patterns of use internal to the speaker but also on what is really out there. (I am alluding, of course, to the famous thought experiment in Putnam 1975.) And yet we might still think that there has to be something in the speaker’s mind, viz., the English word “water” or its Mentalese synonym, that carries that content.

This might be a reason to suppose that all thought is explicit if we had to suppose that language-like representations are the only vehicles of cognition. But, though that is a common prejudice, there is good reason to deny it. It is a remarkable fact that philosophers and psychologists have all but ignored the mind's capacity to use mental imagery as a means of solving problems. The so-called mental imagery debate (e.g., between Pylyshyn 2002 and Kosslyn et al. 2006) has focused on the question whether mental imagery even exists. Once we get past that (the answer being, yes, it exists), a further question is what the mind can do with it. What has never been carefully studied is our capacity to solve problems by mentally picturing *how things go*. We can picture a ball falling, a hammer striking, a person diving from a cliff into the ocean, a hook on a string sinking into the water and a fish biting on it. We can distinguish between imaginings that represent the way things actually go and those that represent what would never happen – a man flying, a tree melting. By means of such mental movies we can plan a course of action and satisfy our needs and desires. What we should expect of conceptual thought is only that it adds something to this process. It need not take over and become the sole medium of cognition.

Precisely because what we can require of conceptual thought is only that it add something to the processes of imagistic cognition and not that it supplant imagistic cognition, conceptual thoughts may remain inexplicit. The reason to invoke mental imagery is not that mental images can be added to an inexplicit thought to make it explicit.⁷ The lobster man's wife's thought does not become a thought about lobsters by having a picture of a lobster appended to it where the word "lobster" should go. A picture of a lobster does not as such represent lobster-kind in general, but a representation of

lobster-kind in general is what would be needed to turn her inexplicit thought into an explicit thought having the propositional content of “He caught more than 30 lobsters today”. Rather, an inexplicit thought can do real cognitive work in the context of a largely imagistic thought process, because even an inexplicit thought may suffice to channel imagistic thought in one direction rather than another.

Mrs. Winkelvoss’s chauffeur faces only two choices, to remain parked or to drive to the entrance, which he may represent imagistically as elaborations on his current perceptual representation. To initiate his drive to the entrance, a cue from the doorman in the form of the incomplete sentence, “Mrs. Winkelvoss is ready”, is all that he requires. The addition of “to go shopping” would draw no further distinction that bears on his choice. Similarly, when the chauffeur thinks to himself, “Mrs. Winkelvoss is ready”, he does not have to add in thought a specification of what she is ready for. If I am pouring oats into a bucket, then at each moment in time as the oats flow, I face a choice between two options, both of which I may represent imagistically as elaborations on my current perceptual representation — continuing to pour or holding back. A cue in the form of the inexplicit sentence, “That’s enough”, is quite enough to decide me between these two options. Similarly, my thinking to myself “That’s enough”, can have the same effect without my having to add in a thought a specification of what I have done enough of (pouring oats into this bucket, as opposed to filling buckets).

A more artificial example might highlight the general principle more clearly. Imagine a subject in a psychology experiment who is presented with pairs of cards containing colored shapes, one per card – a red square, a green triangle, etc. The subject has to point to one card or the other depending on the utterances of the experimenter. If

the two cards always differ in either color or shape, then the experimenter will be able to direct the subject's choice by naming only a color or a shape; the experimenter need not name both color and shape. In much the same way, an inexplicit thought may suffice to drive behavior where the agent faces only a limited choice between imagistically represented options.

These explanations of how inexplicit thought can drive behavior rest on the assumption that inexplicit conceptual thoughts take place in the context of imagistic representations of a limited array of options. For example, I said that Mrs. Winkelvoss's chauffeur faces only two choices, to remain parked or to drive to the entrance. These limits call for some explanation, because there are always other options that could in principle be considered. There are many things Buddy might consider (leaving for a cup of coffee) and many more that he might consider if certain eventualities came to pass (a gun battle in the street). My supposition is not that an agent vividly imagines several complete courses of action and that his choices are strictly confined to these. But I do suppose that there will be a range of actions that the agent is immediately prepared for. The boundary between what is included and what is not need not be sharply defined. Moreover, this range of alternatives may be imagistically represented without the agent's being consciously aware of them all.⁸

Does the point I am making really depend on the hypothesis that there is nonconceptual, imagistic thought? Suppose that all of our representations were discursive, sentence-like representations. We could still have a number of representations that represent our various options; so we could still maintain that inexplicit thoughts can drive a choice between limited options. The answer is that in that case, explicitism could

perhaps be defended by a careful individuation of thoughts. Somewhere in the chauffeur's mind there would be a discursive representation of the activity that Mrs. Winkelvoss will engage in as a result of her being driven by Buddy. So the explicitist might argue that, properly individuated, the thought that Buddy thinks in response to the doorman's utterance includes that representation of her activity, so that the word-for-word expression of his thought would be closer to "Mrs. Winkelvoss is ready *to go shopping*". But if the background sort of cognition is imagistic, then there can be no guarantee that such a gerrymandering strategy will work.

In Wittgenstein's *Philosophical Investigations*, in the course of his musings about the language games of the primitive builders, he at one point raises the question whether "Slab!" is a *shortened* form of "Bring me a slab!" His answer seems to be that, while we can say that, it only means that "Bring me a slab!" plays the role in our language game that "Slab!" plays in theirs. It does not mean that "Bring me a slab!" must be what a builder who calls out "Slab!" really has in mind (*PI*, §§19–20). But someone might resist this conclusion by maintaining that the builder really does have in mind that the assistant should bring him a slab, that "Bring me a slab!" is a correct expression in our language of the thought that the builder has in mind, because when we express in our language what a speaker has in mind, we do so by translating into our language the thought he has in mind, and the translation into our language of the thought that the builder has in mind is precisely this: "Bring me a slab!"

Similarly, one might object to all of my examples of inexplicit thought on the grounds that the thoughts I report as inexplicit are, properly translated, revealed to be explicit. They look inexplicit only if translated into English as if they were sentences of

the communal language. Rather, they have to be taken as sentences of local, special-purpose languages, and, as such, they may be regarded as fully explicit. When the doorman calls Buddy and says, “Mrs. Winkelvoss is ready”, he is speaking a variant of English in which “ready” *literally means* “ready to go shopping”.

If we defend explicitism in this way, then, while we might preserve the letter of the doctrine, we will destroy its animating spirit. Explicitism will no longer support a distinction between spoken language and the language of conceptual thought, because we will no longer have to find a difference between the two languages due to the fact that the sentences in the language of thought have to be explicit in a way that spoken sentences need not be. If we say that the purportedly inexplicit thoughts are really explicit, because they belong to a special language, then we can say the same about the purportedly inexplicit spoken sentences. Accordingly, explicitism will no longer support the common model of linguistic communication that I mentioned at the start.

6. Explanation in terms of inexplicit thoughts

I have argued that inexplicit thoughts can drive behavior because they take place in the context of an imagistic representation of options. But that conclusion does not quite address the second argument for explicitism, because when we explain a person’s behavior, this imagistic situation is not something we typically describe. We do not describe the range of options the agent imagistically represents. Moreover, the agent’s imagistic representation of a range of options is not something we typically have any knowledge of. So, though an inexplicit thought might drive behavior, our explanations

may not be able to cite only the inexplicit thought as such. Is that not a reason to think that thoughts themselves must be explicit?

We may usefully distinguish three cases. One case is the case in which our explanation takes place in much the same setting as the action to be explained. In that case, what we have to explain may be limited in much the same way that the agent's choices are limited. If our explanation of Buddy's action occurs in the same setting as his action, then it may be clear to us, as it is to him, that his choices are limited to remaining parked or driving to the entrance. In that case, we may explain his driving to the entrance perfectly well by saying simply, "He thinks that Mrs. Winkelvoss is ready". A second case is that in which, though we are in a setting removed from the setting of the action, we can describe the setting in a way that makes evident the limitations on the agent's choices. In the case of Buddy, the explanation may begin with an account of the routine that Buddy, the doorman and Mrs. Winkelvoss are engaged in. Given that, we may attribute to Buddy the pertinent thought by saying only "Buddy thought that Mrs. Winkelvoss was ready". It may be clear enough from the prior account of the setting what he thinks she is ready for. Finally, instead of describing the setting, we may make more of the content of the thought explicit in the "that"-clause by means of which we ascribe the thought. In explaining Buddy's actions to the police after the street fight, we might have to say something like, "Buddy thought that Mrs. Winkelvoss was ready to go shopping".

In general, an explanation of behavior that cites the agent's thought will refer to the thought by its content rather than by the structure of the concrete particular representation (the vehicle) that bears that content. More or less of that content may

usefully be mentioned, depending on how much of the setting is independently clear to the interlocutors to whom the explanation is provided. If the setting in which the explanation is offered is removed from the setting in which the action takes place, so that hearers cannot observe the setting for themselves, then to that extent it may be useful to provide more of the content of the thought in the sentence that attributes the thought. In particular, our attributions of thought in explanation of behavior may have to make more of the content of the thought explicit in the “that”-clause than is explicitly born by the vehicle of thought itself.

According to the computational theory of mind, all cognition may be modeled as the rule-governed manipulation of symbolic representations. Perhaps the computational theory of mind would encourage explicitism. Insofar as it ignores the possibility of imagistic thought, it might lead to explicitism in the way I described in the previous section when I considered whether my account of the ways in which inexplicit thoughts might drive behavior really depended on the assumption that there are nonconceptual, imagistic modes of thought. In any case, the computational theory of mind now strikes most students of the mind as a much too limited model of cognition.

According to a certain conception of rational choice, rational choice is a kind of logically valid inference from premises to a practical conclusion, which says, in effect, “Do that!” This conception of rational choice too might lend some support to explicitism. But however well this conception of rational choice might fare as a normative account of what is rationally correct, as a psychological account of actual human decision making it patently fares quite poorly; so I think we can ignore it.

7. The semantics of context-relativity

The explicitist assumes that thoughts must be explicit because there is no place else to look for what determines the propositional content of a thought other than in what we find in the structure of the thought itself. So if we say, “Buddy thinks Mrs. Winkelvoss is ready to go shopping”, then that is true only because Buddy’s thought explicitly specifies that it is shopping that Mrs. Winkelvoss is ready for. If this assumption is to be rejected, then we will need some other account of what makes it the case that a given thought has the content it has. What, if not the semantic contents of the elements that compose an explicit thought, will determine the content of the thinker’s thought?

I will answer this question by explaining what determines the semantic content of an overt utterance. The answer I give will *not* be that it is determined by the content of an underlying thought in the speaker. Consequently, there will be no threat of regress in assuming that the same account might apply to thoughts. The account will fall into two major parts. The first part will be an account of the form of a semantic theory for a language exhibiting context-relativity. The second part will be an answer to a question posed by the first part.

We may assume that a semantic theory for a language permitting context relativity may be specified by clauses of the following form:

(E) Expression e has value v relative to context c if and only if . . . e . . . c . . .

For example, we might have clauses like these:

(D) The word “that” refers to object o in context c if and only if o is the object that c assigns to “that”.

(Q) A sentence of the form, “Every F is G ” is true in context c if and only if for each object o in the *domain of discourse for c* , if o belongs to the extension of F in c , then o belongs to the extension of G in c .

One thing to notice about these clauses is that they determine the values of *expressions* relative to contexts. For example, (D) tells us about the *word* “that” and (Q) tells us about *sentences* of a particular form. They do not assign values to tokens. Also, the right-hand side of the biconditional does not discharge all reference to the expression or to the context. Also, the truth of a sentence can depend on the context relative to which it is evaluated in ways that go beyond what the context assigns to elements of the sentence. As (Q) is formulated, the context-relative truth of a quantified sentence depends on a contextually determined domain of discourse, but that domain need not be *assigned* to the quantifier.

Contexts, as they are spoken of in a theory of this kind, are formal structures of a certain kind. A context will *assign* an object to a demonstrative expression. It will *contain* a domain of discourse. It will contain various other constructs. I will assume that contexts incorporate enough structure that sentence truth may be relativized to contexts alone (and need not be relativized in addition to such things as possible worlds and times). Our semantic theory will evaluate every expression of the language with respect to every context. Contexts, so defined, have to be distinguished from *situations*, considered as an arrangement of concrete events and objects. An *utterance* will be a component of a situation, not a context.

As I have documented elsewhere (2003, 2012a,b), the inexplicitness of sentences does not prevent their standing in logical relations to other sentences. For instance,

“Buddy is ready” presumably implies “Buddy is ready for something”, but not conversely. The utility of the kind of semantics I have described here is that it allows us to define a property of logical validity for arguments composed of inexplicit sentences. An argument is logically valid if and only if for every context in which the premises are true the conclusion is true as well. This virtue provides one of the reasons to distinguish between contexts and situations. The set of contexts can be precisely defined and in that way becomes formally tractable. That an argument is logically valid, or not, will be something we can rigorously prove. Whereas if we said that logically valid arguments preserve truth in all situations or all possible situations, then logical validity would not in the same way be subject to demonstration.

This sort of semantic theory will not all by itself allow us to assign a truth value to a particular utterance, a token of a natural language sentence. For that we need to know in addition which context – which of the formal structures that contexts are – is the one that, as we may put it, *pertains* to that particular utterance. Then we may say:

(U) An utterance *u* of a sentence *s* is true (simpliciter) if and only if *s* is true relative to the context that *pertains* to *u*.

In this way, our account of the conditions under which a sentence is true in a context generates an account of the conditions under which an utterance of a sentence is true simpliciter. We may then identify the *proposition* that an utterance expresses with an explicit statement of the conditions that the world would have to meet in order for the sentence uttered to be true relative to the context that pertains to the utterance.

In light of principle (U), we face the question: What does it take for a context to pertain to an utterance? At this point we might get an argument for explicitism. It might

be said that what determines the context pertinent to an utterance is only the content of an underlying thought. So, for instance, we might say that an utterance of “that” refers to an object *o* if and only if *o* is what is referred to in the thought that underlies the utterance in some sense of “underlie”. If we said that, then we could not, on pain of infinite regress, suppose that the semantics of mental representations was context relative in the same way.

It is not at all obvious that the only possible account of the pertaining relation is one that treats it as determined by the content of an underlying thought. Some features of the context pertinent to an utterance may be determined by the shared goals that the interlocutors share in the conversation in which the utterance occurs together with the contents of the situation that it is part of. For example, some of the elements of a context are domains of various kinds. For the evaluation of quantified sentences, the context needs to contain a domain of discourse (see my 2003). For the evaluation of modal and conditional sentences, the context needs to contain (or *be*) a set of prospects, or “possibilities” (or other contexts) (see my 2005). These domains may be determined by what is objectively relevant to the achievement of the goals of the conversation. The domain of discourse for ordinary quantifiers, for example, may consist of that subset of objects that are actually part of the situation in which the utterance occurs and that have a bearing on the achievement of the goals of the conversation. The domain of prospects relative to which we evaluate a conditional sentence may comprise all and only those that are in some sense live prospects that interlocutors have to prepare themselves for.

Other elements of the pertinent context may be determined by a variety of “accessibility criteria”. These may include such considerations as what is literally salient

in the environment, what the speaker would charitably be interpreted as saying, and how the speaker has used conventional devices such as pointing. For example, the referent of a bare demonstrative “that” may be identifiable on the basis of such criteria. The various criteria may suggest various solutions, and in that case, we may just have to say that the context is whatever reasonable interpreters would take it to be on the basis of an all-things considered judgment. (For a fuller discussion, see my 2008.)

On such an account of the content of the context that pertains to an utterance, the utterer is not authoritative over the meaning of his or her words. The speaker may have one proposition in mind (by virtue of a representation distinct from that which he or she speaks) and his or her utterance may express some other proposition. I find nothing surprising or objectionable in that result. It means that speakers have to take care that their words mean what they intend them to mean. Slightly more surprising is the conclusion that a thinker may misunderstand his or her own thought. The proposition that it bears, in light of the what is objectively relevant and in light of the accessibility criteria, might not be just what the thinker himself or herself would have had said if asked to elaborate. But I think we can tolerate that conclusion as well.

8. A note on subsentential speech

Sometimes we manage to communicate by speaking phrases that are not complete sentences. Sanjay and Silvia may be loading up a moving van. They are looking for a missing table leg. Sanjay says, “On the stoop” and thereby communicates to Silvia the proposition that the missing table leg is on the stoop. This example comes from Stainton (2006) who argues, on the basis of this fact, that natural languages cannot be the medium

of conceptual thought. “Sentence_{semantic}” is Stainton’s term for sentences that express a complete proposition explicitly in my sense (2006, p. 30). The core of his argument is this:

If thought really were inner speech, then occurrently grasping *any* content *c* would require tokening, somewhere in the mind, a natural language expression with content *c*. So as a special case, occurrently grasping a *proposition*, say *p*, would require having a natural language sentence_{semantic} with content *p* running through the mind. (Stainton 2006, pp. 177–78).

Why? Why could the proposition that *p* not be the content of a subsentential *thought*? If, as Stainton maintains, we can make assertions and communicate by means of subsentential expressions, why can we not also *think* by means of them? As far as I can see, everything I said in section 5 applies here as well.

The case of subsentential speech does pose an additional problem for what I have said, however. That is because what I said about the semantics of inexplicit sentences does not carry over without modification to the case of subsentential expressions. We can suppose that for every kind of context-relativity affecting sentences, there is an element of the context that provides an appropriate value. That will make complete contexts very large and complex entities. I do not see any problem in that fact. (A context may leave most parameters simply unevaluated.) However, we cannot suppose that contexts also supply completions of some kind to expressions that are less than complete sentences. That really would make them indefinitely complex.

In reply, one could argue that, although the truth value of a subsentential linguistic utterance is not determined by the semantics of the expression used and the *context*, in my sense, that pertains to the utterance, nonetheless a truth value can be

determined for a subsentential expression by the semantic properties of the expression used together with the features of the *situation* in which it occurs. What we would lose in so appealing to situation in place of context in this way is the license to apply the definition of logical consequence that I presented in the previous section. That seems acceptable. While subsentential expressions may be used to make utterances that express propositions that stand in logical relations to other propositions, subsentential expressions cannot serve as premises or conclusions in logically valid arguments.⁹

References

- Clark, Andy, and David Chalmers (1998). "The Extended Mind", *Analysis* 58: 7–19.
- Corazza, Eros (2011). "Unenriched Subsentsential Expression", *Philosophy and Phenomenological Research* 83: 560–582.
- Gauker, Christopher (1997). "Domain of Discourse," *Mind* 106: 1-32.
- Gauker, Christopher (2003). *Words without Meaning*. Cambridge, MA: MIT Press.
- Gauker, Christopher (2005). *Conditionals in Context*. Cambridge, MA: MIT Press.
- Gauker, Christopher (2008). "Zero Tolerance for Pragmatics," *Synthese* 165: 359–371.
- Gauker, Christopher (2010). "Global Domains versus Hidden Indexicals", *Journal of Semantics* 27: 243-270.
- Gauker, Christopher (2011). *Words and Images: An Essay on the Origin of Ideas*. Oxford: Oxford University Press.
- Gauker, Christopher (2012a). "What Tipper is Ready For: A Semantics for Incomplete Predicates", *Noûs* 46: 61–85.

- Gauker, Christopher (2012b). "Semantics and Pragmatics", in Delia Graff Fara and Gillian Russell (eds.), *Routledge Companion to the Philosophy of Language*. New York: Routledge, 18–28.
- Kosslyn, Stephen M., William L. Thompson, and Giorgio Ganis (2006). *The Case for Mental Imagery*. Oxford: Oxford University Press.
- Merchant, Jason (forthcoming). "Ellipsis", in Artemis Alexiadou and Tibor Kiss (eds.), *Syntax: An International Handbook, 2nd edition*. Berlin: De Gruyter Mouton.
- Perry, John (1986). "Thought without Representation", *Proceedings of the Aristotelian Society, Supplementary Volumes* 60: 137-151.
- Putnam, Hilary (1975). "The Meaning of 'Meaning'", in Keith Gunderson (ed.), *Language, Mind and Knowledge. Minnesota Studies in the Philosophy of Science*, vol. 7. Minneapolis: University of Minnesota Press, 131-193.
- Pylyshyn, Zenon W. (2002). "Mental Imagery: In Search of a Theory", *Behavioral and Brain Sciences* 25: 157–237.
- Stainton, Robert (2006). *Words and Thoughts: Subsentences, Ellipsis, and the Philosophy of Language*. Oxford: Oxford University Press.
- Stanley, Jason (2007). *Language in Context*. Oxford: Oxford University Press.
- Vicente, Agustín, and Fernando Martínez-Manrique (2005). "Semantic Underdetermination and the Cognitive Use of Language", *Mind and Language* 20: 537–558.
- Vicente, Agustín, and Fernando Martínez-Manrique (2008). "Thought, Language and the Argument from Explicitness", *Metaphilosophy* 39: 381–401.

Vicente, Augustín (2010). “Context Dependency in Thought”, in François Recanati, Isidora Stojanovic, and Neftali Villanueva (eds.), *Context-Dependence, Perspective and Relativity*. Berlin: De Gruyter Mouton, 69-89.

Wittgenstein, Ludwig (1953). *Philosophical Investigations*, trans. G. E. M. Anscombe. New York: Macmillan.

¹ The explicitists’ indulgence of demonstratives in thought might lead to an indulgence of inexplicit quantifiers if the explicitist takes on Jason Stanley’s thesis (2007) that quantifiers carry a hidden indexical in their deep structure. I will ignore this, though, since I do not think Stanley’s thesis has won wide acceptance, and, in any case, I have elsewhere argued that it is wrong (Gauker 2010).

² An argument along these lines can be found in works by Vicente and Martínez-Manrique. See Vicente 2010, p. 72, Vicente and Martínez-Manrique 2005, p. 545, and Vicente and Martínez-Manrique 2008, p. 386. However, in their version, the threat of regress arises in the interpretation of the interpreter’s thought, not in the interpretation of the speaker’s thought.

³ Vicente employs such an argument as well. In his 2010, pp. 75–76, he replies to an objection to an argument of the first sort. The reply takes the form of arguing that “relativized truth-conditions” do not suffice for explanations of behavior.

⁴ Vicente, 2010, p. 83, in the context of a discussion of John Perry’s (1986) somewhat similar story of the Z-landers, briefly considers such an objection.

⁵ Vicente, 2010, p. 81, also considers such a response in connection with Perry's example of the Z-landers.

⁶ This is pointed out in Corazza 2011. Corazza thinks that for the supposition that speakers think explicit thoughts he must substitute the assumption that speakers think thoughts about what he calls "reflexive truth conditions", which make reference to the very sentence that bears them. I find it implausible that people, including young children, understand sentences only by explicitly thinking about the truth conditions of sentences. Moreover, as Elugardo demonstrates (this volume), Corazza is at cross-purposes with himself in appealing to reflexive truth conditions while minimizing free enrichment.

⁷ This is the only way of supplementing thought with imagery that Vicente and Martínez-Manrique contemplate, 2008, p. 394.

⁸ For further exposition of the nature of imagistic cognition, including references to empirical literature, and discussion of the imagistic representation of potential courses of action, see my 2011, chapters 5–7.

⁹ I thank Reinaldo Elugardo, Laurence Goldstein and Jeffrey Francis Pelletier for helpful comments on an earlier draft.