

UNIVERSITY OF FLORIDA
DEPARTMENT OF PHILOSOPHY

Narrow Content and Utterance Meaning

Undergraduate Honors Thesis

Spring 2018

By: Irina Bigoulaeva
Faculty Advisor: Dr. John Biro

Table of Contents

I. Introduction

II. Internalism and Externalism about Mental Content

i. Internalism and Externalism Defined

ii. Social Consumerism of Concepts

iii. Double-Content Theories

iv. Puzzling Pierre and Perspectival Content

v. Privileged Access: A Case for Narrow Content

III. Mental Content and Utterance Meaning

i. The Meaning of Natural Kind Terms

ii. Speaker Meaning and Broad Language Compatible

IV. Conclusion

V. References

I.) Introduction

In this thesis, I will defend a version of the strong internalist account of mental content and its implications for how to think about what determines the meaning of utterances. The topic of meaning has been thoroughly examined by H.P. Grice, who distinguished between natural meaning (meaning_N) and non-natural meaning (meaning_{NN}), the second of which is characterized by conventionality as opposed to being grounded in natural laws, such as smoke's meaning_N fire.¹ It is to the second of these two types that the sort of meaning found in language belongs. In fact, as Grice argues, both sentences and speakers may mean_{NN} things: When a speaker makes an utterance P, we may isolate the *speaker's meaning*, that is, what the speaker intended to convey to their audience by means of their utterance; and also the *linguistic meaning* of P, which is the standard meaning the sentence P has in the speech community.² However, it is a topic of debate as to how exactly these two aspects interact to constitute the meaning of a speaker's utterance on the actual occasion of production. On the one hand, it seems that utterance meaning is governed in an essential way by linguistic meaning. I cannot simply say "It's cold in here" to mean that it's warm in here, or that today is Tuesday.³ But neither can a speaker's communicative intentions and mental states be factored out completely. I may be being sarcastic, or using the expression in a nonliteral way, both of which are factors that depend wholly upon me, the speaker, and make the meaning of my utterance different from its linguistic meaning. However, it is not clear how much each of these aspects contributes to utterance meaning, and which one, if either, overrides the other in the end.

¹ Grice (1957), 384.

² Lycan (2008), 86.

³ Lycan (2008), 86. The example of "It's cold/warm in here" derives originally from Wittgenstein.

More light may be shed on this discussion by considering its close connections to the debate on the nature of mental content. In the philosophy of mind, the content of an intentional state, such as a belief, is considered to be the individuating factor of that state. So belief b_1 : "Grass is green" is different from belief b_2 : "Water is wet" in virtue of the two beliefs' differing contents. One is a belief about water, the other, a belief about grass. But in virtue of what is this the case? In other words, what are the factors that determine the content of a belief?

If we can find the answer to this question, we will have a strong tool to assist us in determining the meaning of utterances, at least those that are informative and declarative in nature. Suppose again that I have the belief that grass is green and wish to report it by simply uttering "Grass is green". The speaker-meaning of this utterance will be that grass is green, since I have uttered it with the intention of conveying that proposition. Examining Grice's analysis, we arrive at the conclusion that speaker meaning depends in an essential way on mental content. Grice defines the notion of speaker meaning as follows: To say that a speaker A meant_{NN} something by an utterance x is equivalent to saying that A uttered x with the intention of inducing a belief B in their audience, and for the audience to form this belief on the basis of recognizing the utterance x as so intended.⁴ In other words, A's goal is to get the audience to believe that A believes something. The content of the belief that I have, as well as the content of my desires and goals and other intentional states, will determine the content of the belief I wish to impart to my audience. Therefore, the content of my intentional states will determine my communicative intentions, that is, my speaker meaning.

⁴ Grice (1957), 384.

But the content of my belief will also affect what words I choose in constructing the utterance. Assuming that I have a firm grip on the rules of my language, I will be inclined on the basis of my belief that grass is green to use the term *ōgrassö* in my utterance and not, for example, *ōsandö*, which has a linguistic meaning completely different from the meaning I wish to convey. Thus the content of my belief will predispose me to use words whose linguistic meaning matches the content of my belief, and provided that I choose correctly, the linguistic meaning of my sentence will indeed match my speaker meaning. But if it can be argued, as I will try to do, that this fortunate correspondence is all there is to the role of linguistic meaning in an utterance's meaning – that is, that linguistic meaning does not play a role in actually constituting the meaning of an utterance – then it will follow that the determining factors of the content of my belief that grass is green will also be the determining factors of the meaning of my utterance.

Such an argument is likely to be successful, as we have already seen cases (sarcasm, etc.) where it is undesirable to say that linguistic meaning overrides speaker meaning. But in order to make it plausible that a kind of meaning that is subject-independent and perhaps even beyond one's awareness cannot overrule a meaning that is wholly dependent on one, we must counter the argument that subject-independent factors in general override subject-dependent factors in determining mental content. In the case of declarative utterances that report beliefs or other mental states, my speaker meaning will be the content of the belief I am trying to express, say, b_1 . From what we have seen in Grice's analysis, this seems indisputable. If, however, in my attempt to report b_1 I accidentally end up uttering a sentence whose linguistic meaning is b_2 , and my hearer interprets my utterance in accordance with this linguistic meaning, then my hearer will attribute the belief b_2 to me. This belief attribution will be wrong, precisely because the hearer took my sentence to

mean strictly what it meant, even though that was not the content I was trying to communicate and was thus not what I meant.

However, if it were the case that belief content essentially depended on factors independent of me, factors which I may not even be aware of, then those factors which, by definition, I *must* be aware of could not play the primary determining role. Clearly, I must always be aware of what I mean. If utterance meaning depended on subject-independent factors, my utterance would as a rule not capture, or even represent, the content of the belief I wished to express. At best, it would merely approximate this content.

Thus, it is clear that the nature of the determining factors of mental content will determine the nature of the factors governing utterance meaning. So the first question we must answer is, what is the nature of mental content?

II.) Internalism and Externalism about Mental Content

Internalism and Externalism Defined

Roughly speaking, there are two positions concerning the factors that determine the content of an intentional state. The first is externalism, which states that mental content is determined (at least in part) by factors external to the subject, and the second is internalism, the theory that mental content is determined by factors internal to the subject. Traditionally, the boundary between the two was drawn in terms of some area of the body such as the skin or the brain, with ‘internal’ designating factors inside the area and ‘external’ designating those outside it. Thus many arguments in favor of externalism are based upon cases in which two subjects are in the same state internally (internal ‘twins’), but whose

external circumstances differ and thus result in a difference in mental content.⁵ However, Katalin Farkas has argued that a distinction based on areas of the body is misleading, and that one can formulate the internalist and externalist positions without it. She advances various arguments to show this. The first involves the skeptical brain-in-a-vat argument, in which I imagine that I am a bodiless brain in a vat that is being fed a global hallucination that perfectly simulates the experience of having a body.⁶ Skeptical conclusions notwithstanding, the simple externalist could argue that in the counterfactual situation in which I am simply a brain, I cannot have thoughts that are about my body, because my body does not exist. The important thing to notice is that even when bodiless, my experience would be exactly the same, and yet the contents of my thoughts are different from what they would be if I had a body.⁷ Similarly, the Cartesian demon argument ó an internalist argument ó rests on the fact that even if we were merely tricked by an evil demon into thinking that there is an external world, everything would seem the same as it would if we were not being tricked, from which the conclusion is drawn that the contents of our thoughts would be the same in both scenarios. From this it is clear, as Farkas concludes, that the difference between internalism and externalism is not whether the content-determining factors lie inside or outside the *body*, but whether they are within or beyond the scope of the subject's introspective awareness ó what she calls the subject's "point of view".⁸ The externalist argues that the thought contents of two subjects with introspectively-identical mental states could nevertheless differ, while the internalist argues that their thought contents must be the same.

⁵ Farkas (2008), 75.

⁶ Farkas (2008), 82.

⁷ Farkas (2008), 82.

⁸ Farkas (2008), 82.

In what follows, I will understand the debate between internalism and externalism as Farkas defines it. This still leaves open the matter of determining which position is stronger, however, as well as that of determining the more specific nature of those factors that determine mental content apart from the general description of being subject-external or subject-internal. In Farkas's view, beliefs and intentional states in general are fully characterized by what she refers to as their narrow content. She defines this term as the "phenomenally-constituted intentional features of thoughts and experiences," a definition which rests on an understanding of intentional states as phenomenal states.⁹ However, it will be enough for the purposes of this paper to rephrase her definition of narrow content as consisting of the intentional features of a mental state that the subject has privileged access to, not commenting on the separate topic of whether or not those intentional features reduce to phenomenal states. When I have a certain belief in my mind, according to Farkas, its narrow content is constituted of all of the intentional qualities I have privileged access to in virtue of the belief's being mine, and none that I do not. With respect to Twin scenarios, the aspects that the twins share is their narrow content, while the differing aspects determined by their respective environments is their broad content. According to Farkas's internalist position, however, narrow content fully characterizes a subject's intentional states, leaving no role for broad content. When I have a belief, the content of this belief is just its narrow content.

This is not an uncontroversial position, however, as different philosophers assign to narrow content different properties and roles, if they indeed acknowledge it as a kind of content at all. Brian Loar argues for a double-content theory involving what he refers to as "social" and "psychological" content, corresponding roughly to broad and narrow content,

⁹ Farkas (2008), 93.

respectively.¹⁰ Other philosophers such as Robert Stalnaker and Hilary Putnam have suggested that broad content is sufficient for all explanatory work and that an account of narrow content, if at all possible, yields nothing of interest. My goal in this section is to defend Farkasø definition of narrow content and its role as the sole determining factor of a mental state. In order to do this, I will sketch some rival accounts of mental content from both internalist and externalist perspectives, and argue that Farkasø account of purely-narrow mental content is preferable in all of these cases.

Social Consumerism of Concepts

On the extreme externalist view, the content of intentional states depends solely on outside factors. These factors can be constituted by either the physical environment or the social environment. In a society where there are no scientists, I cannot have a belief that I will become a rocket scientist. If there is no water on my planet, but instead a different drinkable liquid, then I will never have encountered water and have thus never thought about it, even if the liquid in my environment is macroscopically indistinguishable.¹¹

Such intuitive examples might suggest the more general view that the content of my belief is restricted to those concepts that are available in my community. When I have a belief, the content of my belief is always something that another member of my community can grasp in the very same way and attribute to me, and never something that they cannot. Even members of the community who have an impoverished understanding of a concept are nevertheless held to the standard understanding of that concept and defer to it in their beliefs and belief reports. Gabriel Segal refers to this view as "consumerism", a term which

¹⁰ Loar (1988), 99.

¹¹ This situation is presented in Hilary Putnam's famous Twin Earth thought experiment, to be discussed later.

highlights the fact that individual subjects do not determine the content of their beliefs but instead obtain the constituent concepts already in a ‘prepackaged’ state ó presumably from a group of experts in that community that understand all aspects of those concepts.¹²

A well-known argument for consumerism comes from Tyler Burge.¹³ In this argument, we are asked to consider a certain individual who has a number of beliefs about arthritis, including that he has had it for years in his joints, that his parents had it, that stiffening joints are a symptom of arthritis, and many others. All of these beliefs are attributable to him by means of that-clauses that contain the expression ‘arthritis’, where arthritis is understood exactly as we understand it ó an ailment of the joints. Later, however, the patient begins to experience pains in his thigh and comes to form the belief that his arthritis has spread to his thigh. But when he expresses this belief to a local doctor, the doctor informs him that arthritis, by definition, cannot occur in the thigh. The patient then retracts his belief and asks the doctor what is wrong with his thigh. Next we consider a counterfactual situation in which the patient’s physical history and ailments are exactly the same, but in which his community is such that the concept of arthritis that the doctors and other experts share in fact includes the sort of ailment that is afflicting his thigh. The patient’s interactions with linguistic expressions is assumed to be the same as in the first situation, meaning that he has heard the exact same word forms and the same sentences, and up to the point where he visits the doctor, has heard neither a confirmation nor a denial on the part of an expert that the ailment in his thigh is in fact arthritis. Nevertheless, in this counterfactual situation, the patient will lack all of the beliefs which would have been attributable to him with that-clauses containing the term ‘arthritis’ in the sense in which it

¹² Segal (2000), 62.

¹³ Burge (1979).

was used in the first situation. Thus the content of the subject's beliefs varied depending on context, and so, the argument concludes, belief depends on the social environment.

On closer examination, however, Burge's argument is problematic. It is not clear whether in the first situation we are entitled to attribute to the patient belief contents by means of that-clauses containing our term 'arthritis' which, as stated above, designates the arthritis concept that means only an inflammation of the joints. If it were indeed the case that this was the content of the subject's belief, then it is not clear how he could later come to believe that an ailment that is strictly an inflammation of the joints was now afflicting his thigh. Furthermore, the fact that he accepted the doctor's correction shows that his original concept was not that of the community, and thus that his belief content could not have been determined by the community's concept. From this it follows that in the counterfactual situation as well, it is problematic to assume as Burge does that the content of the patient's belief will be determined by the standard concept of his counterfactual speech community.

To see why this is the case in general, let us consider a related example from Robert Stalnaker. Suppose a young child reports to us his belief that his daddy is a doctor, yet lacks the knowledge that a doctor is, in the accepted understanding of the term, a practitioner who is not a quack.¹⁴ Now suppose that the child's father is indeed a quack, and the child has often observed his father employing dubious methods that don't seem to help his patients but does not make the connection that this disqualifies his father from being a doctor. His belief is therefore false ó his father is not in fact a doctor. And yet, the child certainly thinks that his belief is true; he would not report to us a falsehood about his father's profession. Therefore, if we were to hold the child to the standard of *our* doctor-concept, then we would be forced to the conclusion that the child has contradictory beliefs

¹⁴ The example derives originally from Daniel Dennett and is quoted in Stalnaker (1984).

and is therefore irrational. How, after all, could he otherwise hold a belief that his daddy is a doctor when he frequently observes behavior by his father that is inconsistent with his being a doctor? But while clearly we observe such seemingly irrational beliefs in children, we nevertheless do not consider them irrational, at least not to the same degree as we would an adult who is in full possession of the relevant concepts. In fact, those children usually grow up to be fully-functioning adults and come to have the same concepts we do. Therefore, the mistake that lies at the heart of the child's belief has nothing to do with his reasoning, but rather with his concepts. His concept of a doctor deviates from the standard, thus explaining his sincere and informed belief in something that seems to us to be a falsehood.

The important thing to note here is that the child's belief that his daddy is a doctor is not the same as his mother's belief that her husband is a doctor. To the mother, a fully-integrated member of this society, a doctor is somebody who employs proven healing methods and not suspect and ineffective ones. Therefore she would certainly retract her belief if she observed the things that her child observed. This indicates a difference in belief content between the mother and child.

So, we see that at in many situations it is undesirable to report a certain subject's belief in terms of a public concept, since the subject's actual belief interacts with his other intentional states differently than it would if the belief content were in fact picked out by the public concept. This is what we observe in Burge's arthritis case, as well. We would be making a false attribution to the patient in the first situation if we were to consider his belief to be about arthritis in the standard sense. In the counterfactual case, it just so happens that the patient's concept of arthritis coincides with the standard concept, but we would not be entitled to make the respective belief attribution to him simply on the basis of

his being a member of the speech community; rather, we must first have observed his behavior, particularly his reaction to the judgment of experts, to determine what concept he in fact has in mind. From this it is clear that a pure consumerist externalism is not a tenable theory of mental content.

Double-Content Theories

One response to the failure of consumerism is the addition of a narrow aspect to mental content alongside the broad one, to serve as the missing piece that captures those aspects of a subject's belief that the broad description cannot. Theories that make this kind of move are known as double-content theories, since they acknowledge the necessity of a broad content, but at the same time posit a narrow content which is stated to be explanatorily and psychologically necessary, but not possible to capture by broad concepts. One such theory has been put forward by Brian Loar. He makes the distinction between social content, which is ascribable to the subject in terms of broad, public concepts, and psychological content, which is the individuating factor of beliefs, that in virtue of which beliefs explanatorily interact with each other.¹⁵ To motivate his theory, Loar asks us to consider an example that is borrowed from Saul Kripke's paper, "A Puzzle about Belief". The example features a man named Pierre, a monolingual French speaker who has never left France. He has heard many positive things about the city London (which, in his language, he calls "Londres"), among them that the city is pretty. So Pierre becomes inclined to sincerely state: "Londres est jolie" to express his belief that the city called "Londres" is pretty. Later, Pierre moves to London and learns English. Unfortunately, he resides solely in unattractive parts of town, so he becomes inclined to state in English: "London is not pretty". But if one

¹⁵ Loar (1988), 99.

were to ask him in French or in English: 'Is Londres pretty?', then Pierre would readily affirm this, not having realized that the city he calls 'Londres' in French and the city he calls 'London' in English are in fact the same city.

As Loar argues, our ordinary principles of belief ascription would lead us to attribute to Pierre both the belief that London is pretty and the belief that London is not pretty, where 'London' in both that-clauses is understood in its standard, socially-determined sense. In Kripke's paper, the focus is on the fact that this situation gives rise to a puzzle about belief: Since both belief ascriptions are legitimate, what does Pierre believe about London? But for Loar, this situation gives rise to a problem about mental content. Namely, if one and the same broad belief description is true of Pierre, how can we explain the obvious intuition that Pierre has not one belief, but two – one about the city he calls Londres, and another about the city he calls London?¹⁶

Loar's answer to this is that Pierre's two beliefs are the same with respect to their social content, but different with respect to their psychological content. They will interact with his other beliefs in different ways, leading to different dispositions and attitudes. If, for example, Pierre believed in France that if he were to ever live in Londres then he would be living in the same city as Oscar Wilde, then acquired the English belief that he now lived in London, then he would not draw from this the conclusion that he was now living in the same city as Oscar Wilde.¹⁷ As Loar argues, this is not merely a matter of a difference of the two beliefs' psychological roles, but of a difference in something that ought to be regarded as a content, since this difference ultimately lies in the realm of semantics – that

¹⁶ Loar (1988), 103.

¹⁷ Loar (1988), 103.

is, in what Pierre conceives the semantic referent of each term to be.¹⁸ Loar claims that in all cases where two beliefs have the same broad description but nevertheless behave differently psychologically, they ought to be regarded as having distinct psychological contents. Indeed, this seems to be the only route available for us to take, since if we view Pierre's beliefs merely through the lens of broad description, they will seem contradictory. In other words, the core claim of Loar's theory is that if broad content is necessary in determining mental content, then narrow content must nevertheless be posited alongside it in order to capture certain things that the broad description will inevitably leave out or misdescribe. All beliefs, therefore, have two contents: a broad, social content, and a narrow, psychological content, and it is only by taking these two aspects together that we can fully determine the content of a belief.

In the end, however, this theory is vulnerable to a criticism from externalists, namely in that it rests on a problematic assumption made by Saul Kripke. In the following section I will explain Kripke's assumption and its consequences for Loar, and as well as for all double-factor theories that attempt to include broad content alongside narrow content.

Puzzling Pierre and Perspectival Content

As mentioned above, the goal of Kripke's paper about Pierre is to prove that Pierre's situation gives rise to a puzzle about belief. Given that Pierre is inclined to assent both to the proposition that London is pretty and to the proposition that London is not pretty, the question arises as to what Pierre in fact believes about London. We have also seen that Loar's account of narrow content depends on this puzzle existing, namely on the fact that we can indeed report Pierre's beliefs as being the belief that London is pretty and the belief

¹⁸ Loar (1988), 103.

that London is not pretty. As John Biro argues, however, the so-called puzzle about belief is not a puzzle at all.¹⁹ The problem is with Kripke's reliance on what he calls the Principle of Disquotation, which he articulates as follows: "If a normal English speaker, on reflection, sincerely assents to $\neg p$ then he believes that p ."²⁰ Applying the Principle of Disquotation to Pierre's case, we have that since Pierre sincerely assents to both "Londres est jolie" and "London is not pretty", he has two corresponding beliefs: the belief that London is pretty, and the belief that London is not pretty. However, at a closer look, this turns out to be mistaken. Pierre does not in fact have a belief that London is pretty – at least, if the word "London" in the that-clause is taken to designate what we, the third-person observers to the situation, understand by it: namely, the city that Pierre has both heard about in France and also currently resides in. It is clear that Pierre understands neither "Londres" nor "London" in this way. Instead, Pierre has a concept of a pretty city with the name "Londres", and of an ugly city by the name "London", and he does not make the connection that the two cities are the same. So, in reply to the question about what Pierre believes about London, the answer is: nothing. Pierre has no belief about a city that is both the city he has heard good things about in France and the city he takes to be ugly having lived there. Therefore, the that-clauses provided by disquotation do not in fact express the contents of the beliefs that Pierre holds, and so we are not entitled to use disquotation in reporting Pierre's beliefs. In order to be more accurate, as Biro argues, we must use a perspectivized that-clause, in other words, a that-clause whose content more closely matches things as they seem to Pierre.²¹ For example, we may say Pierre believes that the city his friends and family told him about is pretty, and that the city he lives in now is ugly.

¹⁹ Biro, (1984).

²⁰ Kripke (1979).

²¹ Biro, (1984).

The perspectival contents expressed by these that-clauses, in turn, are what Pierre has in mind when he believes *‘Londres est jolie’*, and *‘London is not pretty’*.

Biro’s result that the puzzle about Pierre is not a puzzle after all weakens the motivation for Loar’s psychological content. To be sure, it does indeed show that a notion of perspectivization is necessary in order to properly report a subject’s belief, and that the subject’s mental content would be mischaracterized if one attempted to report it using a that-clause containing concepts the subject doesn’t have. But what it also shows is that it is not necessarily a failure of broad content as such if a that-clause fails to capture the subject’s actual belief content. All that may be needed is a that-clause that makes some additional circumlocutions in order to accurately fix on it.²² And if this really is all that we need to describe the content of each of Pierre’s beliefs, then it is not clear why a double-content theory should be preferable to pure externalism. Indeed, as Biro argues, that-clauses are the only tool we possess for describing the content of a belief.²³ In Pierre’s case, we can say that he believes that the city he has heard about is pretty, while the city that he lives in now is ugly. In the seemingly-problematic case of the doctor’s child, where we cannot say that the child believes that his father is a doctor, we can perhaps report instead that the child believes that his father is someone who tries to cure sick people, as this is a more likely description of what the child has in mind than the fully-loaded concept of *‘doctor’*. The important thing to note is that in all cases, the approximations that take into account each subject’s own understanding of things are nevertheless achieved by using terms with broad semantics. So, since we can get this work done with broad content alone, it seems there is no need to posit narrow content to do it.

²² Biro (1992), 283-290.

²³ Biro (1990), 104.

This weak point of Loar's theory is due to his characterization of broad content as a determining factor of mental content, albeit an insufficient one, and his motivating an account of narrow content on the grounds of this putative insufficiency. Thus his theory, and other theories of this type, fall victim to the externalist counterargument that there is no insufficiency and that belief ascriptions made in broad language are in fact an adequate tool for describing all aspects of the content of a belief. The supplementary narrow content in the double-factor theory is shown by the externalist to not be needed, and so it turns out that the double content theorist has no reason to posit narrow content at all.

This concession to the externalist does not, however, imply that broad content alone can constitute belief content. In the next section I will argue that, while broad content and broad belief ascriptions play a necessary role in belief attribution, what in fact determines the contents of beliefs is a kind of content that is entirely subject-dependent.

Privileged Access: A Case for Narrow Content

We have seen that the externalist makes a valid point by remarking that that-clauses are the only tool we have at our disposal in order to express the contents of our beliefs and other intentional states. If the internalist motivates an account of narrow content on the basis of a failure of a disquotational that-clause in describing the content of a belief, then the resulting account will be flawed. But the defender of a different, more radical notion of narrow content can point out that one should not confuse the function of describing mental content with determining or constituting it. We may very well have no other choice than to use that-clauses with socially-determined contents to describe the contents of our own or others' beliefs, but this does not mean that socially-determined contents in general constitute the contents of those beliefs. A philosopher who argues for such a view is Farkas.

Taking a more radical stance than Loar, she argues that broad content does not determine mental content at all, and that narrow content is all there is to the content of an intentional state.

To recall, Farkas defines narrow content as those intentional features of a state that the subject has privileged access to. Therefore, in order to argue against her radically internalist view and in favor of at least a partial externalism, the externalist must be able to argue not just for the possibility of some broad that-clause ascription describing a subject's mental content, but that broad the concepts in those that-clauses actually constitute that content. In this section I will argue that this is possible, but only if the content of the belief ascription is a content the subject has already accessed through his first-person knowledge. In other words, the accuracy of a broad ascription depends on the subject's conceiving his own intentional state in a way identical to the content of the ascription, from which it will follow that a belief's having a certain broad content presupposes its having a certain narrow content. The core motivation for this argument will be that externalism and privileged access are incompatible.

The problem of privileged access has been shown to be a nontrivial one for the externalist. It is generally accepted that we seem to have a privileged kind of knowledge of our mental states simply in virtue of those states' being ours. As Loar argues, if I have a belief that Freud is in Vienna, I know that my belief is about Freud, seemingly without having to have knowledge of any kind of external relation between myself and Freud.²⁴ If we accept that mental content is essentially determined by external relations, however, then it seems we have to either deny that we have any special, authoritative knowledge of the contents of our own mental states or explain this authoritative knowledge in external terms.

²⁴ Loar (1987), 100.

Many have been reluctant to concede the first point and argue instead for the second. Stalnaker is one who attempts to take the second route, namely, to argue that externalism and privileged access are compatible.²⁵ He constructs a thought experiment of a machine that is intended to function as a simplified analogy of a normal, rational subject with beliefs, his argument being that if the machine can form something like reflective beliefs with externally-determined contents, then so can a human being. The machine is designed to perceive passing figures on a conveyor belt and enter certain states based on the figures' attributes. Importantly, the machine's states can be described in terms of propositional content, just like the belief states of a rational subject. If a blue cube passes by the machine, the machine will enter state B_c , which is about the blue cube; if a yellow sphere passes by, it will enter state Y_s , which is about the yellow sphere, and so on. In addition, the machine also has second-order 'memory' states that are generated from the first-order 'perceptual' states: Its first memory state will be B_{c1} if the first item perceived was a blue cube, and Y_{s1} in the case of a yellow sphere. These second-order memory states are analogous to a subject's reflective beliefs about her intentional states.

Clearly, Stalnaker remarks, the propositional content of the machine's first-order states is externally-determined. The machine's entering state B_c is triggered by a blue cube appearing in front of it, in other words by the perceptual relation existing between the cube and the machine. But the machine also has the ability to form reflective beliefs on the basis of those externally-determined intentional states, namely, by directly deriving the contents of the reflective states from the content of the first-order ones. If the phenomenon of privileged access is simply our forming reflective beliefs about our intentional states in an analogous way, then the contents of these reflective beliefs derive directly from the

²⁵ Stalnaker (1990), 142-143.

contents of our first-order states. And so, Stalnaker concludes, if the contents of our mental states are externally-determined, then so are the contents we are aware of through privileged access to those states. Externalism is no hindrance to the formation of a reflective belief, and therefore no hindrance to a rational subject's ability to know the contents of her own beliefs. This is claimed to hold even in cases of perceptual error. If, due to unusual lighting, the perception machine mistakes a blue cube for a yellow cube, it will enter state Y_c , and the content of the machine's state will be Y_c even though there is a blue cube in front of it.²⁶ Similarly, if we are led to perceive an object in front of us by an illusion, the object we perceive is the content of our perceptual belief, even though we do not in fact stand in a relation to such an object.

This argument, however, leads to an inconsistency. Namely, if the external reference relation between me and the referent of my intentional state wholly determines the content of the intentional state, it is not clear why the machine's external relation to a blue cube could ever result in a reflective state that is about a yellow cube. The only way Stalnaker can allow for the machine's reflective state to be about a yellow cube is if the content of the first-order state is also about a yellow cube, which can be the case only if the machine in fact stands in an external relation to a yellow cube. But there is no yellow cube present in the environment; what makes it seem to the machine that the cube is yellow are physical factors that have nothing to do with the object itself, for example, unusual lighting. But surely the machine's first-order state is not about the blue cube plus the lighting – in fact, a blue cube does not figure in to the machine's perceptual belief at all, nor does the lighting. The machine's first-order state is about a yellow cube. In other words, the state has a content such that the effects of the lighting are already factored in. In fact, we can say

²⁶ Stalnaker (1990), 143.

that the content of the machine's state depends on how the lighting makes the cube appear to it. But this is just to say that the content of the machine's intentional state depends on something other than an external relation, namely, the nature of the machine's grasping of the relation. In terms of the rational human subject, it would be the nature of the subject's conscious grasping of their external relation to the object of their intentional state.

This shows that in order for Stalnaker to be able to claim that the content of a reflective state derives directly from the content of a first-order state, he must concede that the content of the first-order state already takes into account factors other than the externally-determined relation between believer and object. These factors are those conscious effects on the believer that additional object-independent factors have, such as unusual lighting making a blue cube appear yellow. These factors themselves are certainly physical and independent of the believer, just as the object itself is. In fact, the believer may not even be aware that these additional factors are in play, as is the case with many perceptual illusions. But their effects must nevertheless be conscious to the subject, as this is exactly what it means for a blue cube to appear yellow. Had they not been, then the blue cube would have been perceived properly as blue, and the content of the first-order state would have been *Bc*. But if we concede this, then we have conceded that the contents of both a first-order state and that of the reflective state derived from it are determined by factors that the subject has already grasped. In other words, in an attempt to hold on to the notion of privileged access, the externalist ends up having to concede that subject-dependent factors determine mental content.

A defender of Stalnaker may nevertheless counter that the external perception relation between me (or the machine) and the blue cube does exert some influence on the content of my (or the machine's) perceptual belief about a yellow cube. After all, my belief was about

a yellow cube and not about a yellow sphere. However, the difficulty this argument presents is minor. We can easily abstract from the perceptual example and suppose that I am not perceiving, but daydreaming. Suppose that I am standing before a blank wall and imagining what it would be like if there were a unicorn in front of me. I picture the unicorn, and on the basis of this picturing I form the belief that the unicorn's horn is purple. Clearly, there is no unicorn in front of me to have caused this 'perceptual' belief, although a neurologist might point out that my act of daydreaming was in fact my manipulating certain electrical signals in my brain and causing some neurons to fire in order to trick myself into thinking that I am seeing a unicorn. But the content of my belief that the unicorn's horn is purple is clearly not determined by an external relation between me and any sort of object in my environment. The only physical objects that I could conceivably be in a relevant relation with are my neurons, but surely the neurons do not conceptually figure into my belief that the unicorn's horn is purple. One might object again and say that the relation that holds between me and the image I have produced is external. But this would not explain why the moment my conscious perception of my daydream changes (for example, when I go from imagining a purple horn to imagining a golden one), my belief changes as well. In other words, merely citing some external relation leaves out the nature of my awareness of that relation, which in the end is the factor that determines my perceptual belief. My belief is about a golden horn because I perceive it as a golden horn, not because the horn is actually golden (in fact, there is no horn at all). Similarly, with regards to the earlier objection, my perceptual belief is about a cube and not a sphere because my conscious grasping happens to be of a cube, not because there luckily is indeed a cube in front of me. As soon as one attempts to characterize this content-determining factor of the belief as

external, one is confronted with the problematic situation that a mere change in conscious perception of the belief content will change the content of the belief.

And so, we have the result that privileged access is incompatible with the broad character of mental content alone determining the content of an intentional state. To be sure, we can certainly speak of beliefs in terms of their external referents, but what is more relevant to the actual content of those beliefs is the intentional character that the believer has accessed by way of first-person awareness. We have seen that this accessed content is not determined by external factors, although it certainly can and must be expressible in broad language. But despite this, it is also the case that privileged access operates independently of external relations between believer and object. Additionally, with the preceding arguments about perspectivized content in mind, we may say that this accessed content serves as the standard by which we may judge the accuracy of a belief ascription. We have already seen that a that-clause can serve as an accurate belief ascription only if it contains certain circumlocutions in order to take into account the subject's understanding of the concepts that figure into their belief. But unless there is a fact of the matter as to which circumlocutions are better approximations than others, the entire endeavor would be in vain. Now, however, we can give a precise definition of what that fact of the matter is, namely the subject's understanding, which is just the content that the subject has accessed by way of first-person awareness of her beliefs. If the subject is unaware of some aspect of the relation between herself and the object of her belief, then the content she accesses would not include those aspects of the relation, even though, in external terms, those aspects of the relation still hold. If we were to claim otherwise, then we would be arguing that the subject always grasps the contents of her beliefs in their fullest, broad sense. But

clearly this is not the case, as we have seen many cases of ignorant subjects such as the child and Pierre.

This leads us to conclude that that the determining factor of one's mental content is not in the end an external relation, but whether or not a certain intentional feature figures into the subject's first-person understanding of the thought she is thinking. This is exactly Farkas's definition of narrow content. And so we have that narrow aspects play an essential role in determining mental content.

III.) Mental Content and Utterance Meaning

Having established that the content of intentional states is not determined by external factors but by internal factors alone, we now turn to the question of whether the same is true of utterance meaning. As explained in the introduction, if it had been the case that the content of my belief essentially depended on factors that could potentially be beyond my awareness, then it would have followed that speaker meaning could not play a determining role in utterance meaning. But it does not immediately follow from this that if all mental content is narrow, then speaker meaning determines the meaning of our utterances. In this part I will provide the argument needed to show that this is the case.

In Grice's work, we find the argument that speaker meaning is wholly constitutive of utterance meaning. This conclusion is achieved by means of a reduction of linguistic meaning to speaker meaning. As was argued in the previous section, language is the only tool we have at our disposal if we want to report beliefs, where language was understood to be broad. If Grice's reduction is successful, however, then it would follow that the meaning of our belief reports is based entirely on our speaker meaning, which is again based on our narrowly-constituted belief content. It would follow that language is not broad, but narrow

as well. But Grice's procedure is not the only way of making an argument for speaker meaning determining utterance meaning. Specifically, we do not need to maintain that linguistic meaning reduces to speaker meaning in order to argue that speaker meaning determines utterance meaning; it will be enough to show that, in all cases, the linguistic meaning of an utterance simply takes second place and is overridden by speaker meaning. It is this looser version of the claim that I will defend here.

First, however, I will explain why Grice's reduction is problematic.

The Meaning of Natural-Kind Terms

One central topic in the debate about utterance meaning is the meaning of natural-kind terms. Natural-kind terms are terms that are used to designate a class of things in the natural world that have something in common by virtue of which they are a member of that class. The word 'diamond' for example, is a natural-kind term, as it picks out all substances that have a certain molecular structure and excludes those that do not. Surely it seems counterintuitive to suppose that the meaning of the word 'diamond' on my mouth is something wholly dependent on my understanding of the substance I refer to by that name, even if I know the word well enough to know that it designates only stones that are hard, clear, and valuable. In fact, if I am like most people, I will readily acknowledge that my understanding of 'diamond' (or, equivalently, my concept of a diamond) is inferior to that of chemists' and will go to them to determine which one of two stones that seem phenomenally identical to me is a member of the class denoted by 'diamond' and which is not. Importantly, once the experts enlighten me on the chemical properties of diamonds, I would realize that these chemical properties are necessary in order for something to be a diamond and therefore that the word 'diamond' on my mouth and everybody else's actually

means something with *this* chemical structure as opposed to *that* one. Natural-kind terms therefore seem like something borrowed, words whose meanings might not be fully grasped by the subject but nevertheless carry their full weight within the speech community at all times. The meaning of the word ‘diamond’ depends on what it is to be a diamond, which is a fact independent of me. So it seems that, at least for natural-kind terms, an externalist theory of meaning is necessary.

A famous argument for this view can be found in Putnam’s ‘The Meaning of ‘Meaning’’. In this paper, Putnam’s goal is to argue against what he calls the ‘traditional’ view on meaning, according to which terms are treated as ordered pairs of sorts consisting of a sense, or intension, and a reference, or extension. The extension is simply the set of things of which the term is true, while the intension is something like the concept associated with the term. It is noted that expressions with different intensions can have the same extension, as with ‘creature with a heart’ and ‘creature with a kidney’ designating the same animals.²⁷ It is also accepted that grasping the intension of a term is a matter of being in a certain psychological state, where a psychological state is understood as a state that does not presuppose the existence of anything other than the subject who is in that state, not even the subject’s own body. Putnam calls this kind of psychological state a psychological state in the ‘narrow sense’.²⁸ It is easy to see that this definition of a narrow psychological state accommodates Farkas’s account of narrow states, which are limited to those intentional states that are phenomenally accessible to the subject. But additionally, the traditional view of meaning claims that the meaning of a term is exhausted by the term’s intension, which implies that knowing the meaning of a word is just a matter of grasping that intension, that

²⁷ Putnam (1975a), 217.

²⁸ Putnam (1975a), 221.

is, being in a certain narrow state. And so, the traditional view provides us with two main claims: (i) that knowing the meaning of a term is just a matter of knowing its intension, which is being in a certain narrow psychological state, and (ii) that the intension of a term unambiguously determines its extension.²⁹ From these claims it would follow that being in a certain narrow state determines the extension of a term. And so, from the traditional view of meaning, we would be able to argue to the Gricean conclusion that being in a certain narrow state ó perhaps the state of having certain communicative intentions ó determines the extension of a word and thus the meaning of the utterance in which it occurs.

Putnam, however, problematizes such a conclusion. He argues that knowing the meaning of a term is in fact not simply a matter of being in a certain narrow state, which implies that there is more to the meaning of a term than just the term's intension. To show that this is the case, Putnam designs a thought experiment, the goal being to show that it is possible for two speakers to be in the same narrow psychological state and still mean different things by a certain word. In the thought experiment, we are asked to imagine a distant planet exactly like Earth, populated with identical copies of all of Earth's inhabitants. But this planet, called Twin Earth, differs from Earth in one respect ó namely, the liquid that fills its lakes and oceans and is used to quench thirst is not H₂O, but a chemically-different substance, XYZ, though Twin English speakers on Twin Earth refer to this substance by 'water', just as we refer to H₂O by the same word. Further, Putnam asks us to imagine that the year is 1750 and that neither planet is chemically literate. So, no methods currently exist for detecting the physical property of each liquid that distinguishes it from the other. Now we imagine Oscar, an Earthling, and Twin Oscar. Both of them experience in their environment a liquid substance with the same phenomenal properties:

²⁹ Putnam (1975a), 219.

clear, wet, odorless, and drinkable. That is, their experiences are subjectively indistinguishable; from the standpoint of both Oscarsø introspective awareness of their intentional states, their concepts of the two liquids are the same. And yet, as Putnam argues, the meaning of the word *õwaterö* on their respective mouths is different: for Oscar it means H₂O, while for Twin Oscar it means XYZ.

The explanation of such cases is a phenomenon that Putnam calls the division of linguistic labor, which is the idea that within a single linguistic community, it is not necessary for all its members to know the difference between a member of the natural-kind class and a phenomenally-similar non-member.³⁰ Going back to the case of diamonds, it would certainly be of interest to many people whether a piece of jewelry is made of actual diamonds or fake ones. And yet it is not necessary that all of these people know what distinguishes a fake diamond from a real one in order for the word *õdiamondö* on their mouths to mean the real thing ó all that must be done is to rely on the definition given by the subset of experts in their speech community who know the relevant necessary and sufficient conditions for something's being a diamond. Thus, as Putnam argues, the factors that determine of the meaning of the word *õdiamondö* are present not in any individual speaker's mind, but in the linguistic community as a whole, and one is a member of this linguistic community only if one conforms to the accepted definition. Even in the Twin Earth case, where there are as yet no speakers capable of isolating the essential property of the liquid they call *õwaterö* and thus providing a proper definition, such a property is presumed to exist. Furthermore, Putnam suggests that one can think of the division of linguistic labor as operating across time, as well, which means that members of a speech community living at a point in the past can and do defer to the meaning of a natural-kind

³⁰ Putnam (1975a), 227.

term as the community's later members, who eventually managed to discover the natural-kind's essential property, came to understand it.³¹ Both main claims of the traditional theory of meaning have been abandoned: the claim that being in a certain narrow state determines the extension of a term, and that the meaning of a term is exhausted by the narrow state.

In effect, Putnam's Twin Earth argument attempts to establish the existence of a linguistic meaning. To define the extension of a term is just to define the things the term may be correctly used to describe and thus to outline a particular community's convention. By disputing Putnam's conclusion of abandoning the above assumptions, one would be arguing that there is no component of the meaning of a term such that the meaning is dependent on an external factor, which is essentially an argument against a term's having a linguistic meaning at all. But, surely, this is an undesirable result. Firstly, it seems problematic to dispute the conclusion that the term 'water' had a different extension on Earth and Twin Earth in 1750. Indeed, there was only H₂O or XYZ present in each respective environment, and additionally, neither planet knew of the existence of the other. So, on each planet, the set of things the term 'water' denoted was necessarily restricted to specimens of one particular substance, even though the inhabitants on both planets understood the respective substances in the same way. Furthermore, maintaining the view, *contra* Putnam, that a narrow state determines a term's extension makes the proposed theory of meaning susceptible to criticism aimed at Grice. By positing that a narrow state may determine a term's extension and thus wholly constitute the term's meaning, we would be making the argument that a notion of linguistic meaning ultimately reduces to some narrow state, potentially speaker meaning. Several objections to Grice show the

³¹ Putnam (1975a), 229.

implausibility of such a claim. William Lycan, for example, points out that linguistic meaning is necessary in order to serve as a standard for what speakers normally mean by an utterance. He remarks that if linguistic meaning reduced to speaker meaning, then there would be no pre-established fact about how to interpret a novel utterance, especially when the novel utterance in question is meant nonliterally.³² And yet, we manage to successfully interpret the meaning of such sentences every day without a problem. Without linguistic meaning, then, we would lose an important standard with which to compare deviant uses, in order to be able to classify those uses as deviant in the first place.

Speaker Meaning and Broad Language Compatible

Consequently, we must accept Putnam's denial of the claim that narrow states determine extension, and therefore deny that narrow states such as speaker meaning may determine linguistic meaning. In other words, there is indeed an aspect of meaning that figures in an utterance and is determined by external factors. However, there is a different stipulation in Putnam's argument that must be challenged, namely that a term's linguistic meaning is all there is to the meaning of an utterance of the term produced by a speaker. In other words, there is reason to regard a speaker's grasping of the term's linguistic meaning, which is a narrow psychological state, as constituting the meaning of the term on the speaker's mouth in a more immediate and fundamental way than the linguistic meaning of the term does.

To see why, let us consider the case of the doctor's child once more. The child has the narrowly-constituted belief that his daddy is someone who cures sick people and reports this belief with the utterance: "My daddy is a doctor." Agreeing with Putnam, we can

³² Lycan (2008), 94.

certainly speak of the linguistic meaning of the child's utterance and claim correctly that, regardless of the child's internal state, the extension of the term *õdoctorõ* that occurs in his utterance is the set of all and only legitimate doctors. However, there is a further element of the utterance that needs describing, one that defines a property of the utterance that ought to be called a *ñmeaning* in its own right. This element is the narrow state that the child is in, namely, his grasping of the concept *õdoctorõ*. In the preceding part of this paper, we saw that the nature of this grasping in fact fully determines the content of his belief. Therefore, if we wish to use the child's utterance to glean an informative picture of what is actually going on in his mind, our only option is to consider the determining factor of the child's belief to be a determining factor of an alternate, additional kind of meaning that exists alongside the linguistic meaning of the utterance, namely, speaker meaning.

What motivates putting speaker meaning at least on a par with linguistic meaning is the fact that if the speaker meaning of the utterance deviates from the linguistic meaning, then an aspect of the child's utterance called its realization conditions will be correspondingly altered. The notion of a realization condition comes from Loar, who defines it as the set of possible worlds in which the speaker's beliefs would be true if he were not subject to misconceptions.³³ Within the framework of his theory, Loar claims that these realization conditions are what constitute narrow content. As we have previously discussed, however, this notion of narrow content is problematic, and we instead argued for a different understanding of narrow content to be the determining factor of mental content. But I will adapt Loar's notion of realization conditions to a different purpose, viz., to be what is determined by speaker meaning.

³³ Loar (1988), 108.

To see how this works, suppose that the child's conception of what it means to be a doctor were not a misconception, but in fact a proper understanding of some other socially-standard concept. We would then have that the child's utterance, "My daddy is a doctor", is true if and only if his father is someone who matches the concept which the child associates with "doctor". In other words, the realization conditions of the child's utterance would be all the possible worlds in which the father falls within the extension of the hypothetical term doctor_C , where doctor_C has a linguistic meaning equivalent to the child's speaker meaning when he uses the standard term "doctor" in our understanding of it. So, perhaps the linguistic meaning of doctor_C is "someone who tries to cure sick people". From this it would follow that the extension of this term is just the set of all people who try to cure sick people. Again following Putnam, this extension of doctor_C is not determined by the child's narrow belief state; it is determined by what it means to be a doctor_C , which just happens to be grasped by the child in full. But the important thing to note is that it is from these possible worlds — the worlds in which the child's deviant grasping of our term "doctor" is in fact a proper grasping of some different term with a different extension — that we pick out the linguistic meanings in terms of which the child's utterance is to be interpreted, and thus the truth conditions according to which the utterance is to be evaluated. Specifically, we would evaluate the truth of the child's utterance according to the reference of the term doctor_C , not according to the reference of our term "doctor". This is reflected in our common social practices. We clearly do not hold children to the same conceptual standards as adults and are more willing on the whole to say that a child who described a quack as a doctor has made a true statement than we would in the case of an adult who made the same claim. And this is achieved by means of evaluating the truth of the child's utterance according to the truth conditions of the possible world in which the

linguistic meaning of the child's utterance coincided with his speaker meaning. In doing this, we do not dismiss the fact that the child's utterance is false on the basis of the linguistic meaning of the word 'doctor'; we simply concede that in spite of this, the child has nevertheless made a true statement, namely that his father was an example of the kind of person that, in his view, exemplified the meaning of 'doctor'. Importantly, the realization conditions do not determine the truth conditions that govern the utterance in each respective possible world; they simply determine which of those possible worlds we will consider relevant to interpreting the utterance.

But if this case can be generalized, then it will justify our claim that speaker meaning has the final say over linguistic meaning in determining the meaning of a speaker's utterance. For we will have that the realization conditions of the utterance, which are determined by the speaker's narrow states, determine which set of externally-determined truth conditions we ought to hold to be the standard against which to evaluate the utterance. And, indeed, we see that the same reasoning used for the child can be applied to Pierre's case. As we recall from the previous part, Pierre has two beliefs: one about a city he has heard good things about in France and another about a city he now lives in. To report these two beliefs, Pierre produces sentences in French and English whose linguistic meanings are, respectively, that the city he both lives in now and has heard good things about is pretty and that the city he both lives in now and has heard good things about is not pretty. But surely this is not what Pierre's speaker meaning was, and surely it is not the meaning that we would ultimately use as a guide to interpreting Pierre's further behavior. After all, the goal of interpretation is to attribute meanings to Pierre's utterances that will make sense of other things he does or says. If we interpret his utterances solely on the basis of their linguistic meaning, we would be led to the conclusion that he is irrational, which is

unacceptable, as it is clearly not the case. Therefore, in order to determine the actual meaning of Pierre's utterance, we must discover Pierre's speaker meaning, which determines a particular set of realization conditions that provide a fact of the matter as to which concepts Pierre has and thus what he means by the words he uses. Analogously to the case of the child, it will turn out that Pierre is not using our term "London". Supposing his understanding of things is not mistaken, we will have that he is using two terms: "London_F" and "London_E", where the referent of the first is the city he has heard about in France and is located in one particular geographical location, and the referent of the second is the city he lives in now, which is located in a different geographical location. (Obviously, the locations of the two cities he conceptualizes must differ, since he considers the cities to be different.) It just so happens that in the actual world, the referent of "London_F" and the referent of "London_E" is the same, leading us to correctly conclude that Pierre is getting something wrong, just like the child is. Nevertheless, we will evaluate Pierre's statements according to the linguistic meanings (in other words, the denotations) of "London_F" and "London_E", leading to the correct evaluation of Pierre's two utterances as meaning, respectively, that the city he has heard about in France is pretty and that the city he lives in now is ugly.

Note that the interpretation will be made in the same broad language that Pierre speaks (and shares with us), but the terms that figure in to our translational "means that" clause will be different than the ones Pierre used to construct his utterance. The terms in our translational clause will be circumlocutions that yield a linguistic meaning that is equivalent to Pierre's speaker meaning. Analogous to the case of mental content, the fact that establishes the correctness of our circumlocutions is the correspondence of the resultant linguistic meaning of our translation to the subject's speaker meaning. And so,

after hearing Pierre's two utterances, we would correctly translate them to mean that the city he has heard about in France is pretty and that the city he lives in now is ugly, respectively, just as we would correctly take the child's utterance to mean that his daddy is someone who cures sick people. Therefore, without speaker meaning, we would not be able to determine the meaning of an utterance.

Certainly, linguistic meaning often proves necessary in order for us to be able to discover speaker meaning. As Lycan points out, without the knowledge of what a word conventionally means in our speech community, we would have no way to even begin to decide what a string of uttered words meant, nor would we have any reason to prefer one hypothesis over another. Nevertheless, as common practice attests, once speaker meaning is discovered, we readily revise our previous belief attributions and utterance interpretations made solely on the basis of linguistic meaning to accommodate it, as we indeed should if we want to accurately explain and respond to the speaker's behavior.

Thus we have the result that speaker meaning consistently overrules linguistic meaning in interpreting utterances and therefore cannot be excluded from a theory of utterance meaning. Our result is consistent with Putnam's argument about extension determining the meaning of a term, without the further implication that this externally-determined meaning, linguistic meaning, is all there is to the meaning of a speaker's utterances. Thus we arrive at an account of utterance meaning that preserves some Gricean elements but deviates from the Gricean reductionist spirit. On our account, the linguistic meaning of an utterance, while a vital ingredient in communication and a necessary tool for acquiring understanding of other people's utterances, nevertheless does not have the final say with regards to the meaning of a person's utterance, just as ascribable broad content does not determine the content of someone's belief. These jobs are left for speaker meaning and narrow content,

respectively. Being able to affirm both broad and narrow aspects of utterance meaning while separating their functions and roles leads to a theory of meaning that is more complete than one that focuses only on one half.

IV.) Conclusion

In this thesis I have argued for a theory of utterance meaning that assigns a role both to linguistic meaning and speaker meaning, but one in which speaker meaning overrides linguistic meaning in cases of the two diverging. I began with the observation that if mental content was essentially determined by subject-independent factors, then it would be impossible for a subject-dependent meaning to constitute the meaning of an utterance. Next I argued that mental content was in fact determined by subject-dependent factors, implying that all mental content is narrow. Nevertheless, I argued against a complete reduction of linguistic meaning to speaker meaning, arguing from Putnam and from a direct counterexample to Grice that an independent notion of linguistic meaning was necessary. Finally, however, I established that speaker meaning overrides linguistic meaning in an important way in interpreting utterances.

V.) References

- Biro, J. I. 1984. 'What's in a belief?', *Logique et Analyse* 107, pp. 267-82.
- ô 1990. 'Individualism and Interpretation', *Acta Analytica* 6, pp. 99-112.
 - ô 1992. 'In Defence of Social Content', *Philosophical Studies* 67, pp. 277-293.
- Burge, Tyler. 1979. 'Individualism and the mental'. *Midwest Studies in Philosophy* 4 (1):73-122.
- Farkas, Katalin. 2008. *The Subject's Point of View*. Oxford University Press.
- Grice, H. P. 1957. 'Meaning'. *Philosophical Review* 66 (3): 377-388.
- ô 1968. 'Utterer's Meaning, Sentence-Meaning, and Word-Meaning'. *Foundations of Language* 4 (3): 225-242.
 - ô 1969. 'Utterer's meaning and intention'. *Philosophical Review* 78 (2): 147-177.
- Kripke, Saul A. 1979. 'A puzzle about belief'. In A. Margalit (ed.), *Meaning and Use*. Reidel. pp. 239-83.
- Loar, Brian, 1987. 'Subjective Intentionality'. *Philosophical Topics* 15 (1): 89-124.
- ô 1988. 'Social content and psychological content'. In Robert H. Grimm & D. D. Merrill (eds.), *Contents of Thought*. University of Arizona Press.
- Lycan, William G. 2008. *Philosophy of Language: a Contemporary Introduction*. 2nd ed., Routledge.
- Putnam, Hilary. 1975a. 'The meaning of 'meaning''. *Minnesota Studies in the Philosophy of Science* 7: 131-193.
- ô 1975b. *Mind, Language, and Reality*. Cambridge University Press.
- Segal, Gabriel M. A. 2000. *A Slim Book about Narrow Content*. Cambridge: The MIT Press.
- Stalnaker, Robert. 1984. *Inquiry*. Cambridge University Press.
- ô 1990. 'Narrow Content'. In C. Anthony Anderson & Joseph Owens (eds.) *Propositional attitudes: the Role of Content in Logic, Language, and Mind*. Stanford: CSLI/Stanford.