

Advances in Experimental Philosophy

Series Editor:

James Beebe, Associate Professor of Philosophy, University at Buffalo, USA

Editorial Board:

Joshua Knobe, Yale University, USA

Edouard Machery, University of Pittsburgh, USA

Thomas Nadelhoffer, College of Charleston, UK

Eddy Nahmias, Neuroscience Institute at Georgia State University, USA

Jennifer Cole Wright, Psychology, College of Charleston, USA

Joshua Alexander, Siena College, USA

Empirical and experimental philosophy is generating tremendous excitement, producing unexpected results that are challenging traditional philosophical methods. *Advances in Experimental Philosophy* responds to this trend, bringing together some of the most exciting voices in the field to understand the approach and measure its impact in contemporary philosophy. The result is a series that captures past and present developments and anticipates future research directions.

To provide in-depth examinations, each volume links experimental philosophy to a key philosophical area. They provide historical overviews alongside case studies, reviews of current problems and discussions of new directions. For upper-level undergraduates, postgraduates and professionals actively pursuing research in experimental philosophy these are essential resources.

Titles in the series include:

Advances in Experimental Epistemology, edited by James R. Beebe

Advances in Experimental Moral Psychology, edited by Hagop Sarkissian
and Jennifer Cole Wright

Advances in Experimental Philosophy and Philosophical Methodology,
edited by Jennifer Nado

Advances in Experimental Philosophy of Mind, edited by Justin Sytsma

Advances in Religion, Cognitive Science, and Experimental Philosophy,
edited by Helen De Cruz and Ryan Nichols

Advances in Experimental Philosophy of Language

Edited by
Jussi Haukioja

Advances in Experimental Philosophy

Bloomsbury Academic

An imprint of Bloomsbury Publishing Plc

B L O O M S B U R Y
LONDON • NEW DELHI • NEW YORK • SYDNEY

- Nichols, S., Pinillos, A. and Mallon, R. (forthcoming) (NPM), 'Ambiguous reference'. *Mind*.
- Nisbett, R. E., Peng, K., Choi, I. and Norenzayan A. (2001), 'Culture and systems of thought: holistic vs. analytic cognition'. *Psychological Review*, 108, 291–310.
- Putnam, H. (1973), 'Meaning and reference'. *The Journal of Philosophy*, 70, 699–711.
- Reimer, M. (2010), 'Reference'. *The Stanford Encyclopedia of Philosophy*. Ed. E. N. Zalta, <<http://plato.stanford.edu/archives/spr2010/entries/reference/>>.

Testing Transparent Ascriptions: A Plea for an Experimental Approach

Mark Phelan

Consider the sentence, "I believe that your keys are in the car". It seems natural to interpret this sentence as saying something about my psychological state, namely, about my beliefs. But now suppose you are desperately searching for your keys, and I tell you, "I believe that your keys are in the car". In this context, this phrase would be typically understood to primarily convey something about the world – about a possible location of the keys. But how can it be that the exact same phrase conveys such different information – about my psychological state, on the one hand, and the external world, on the other?

I will call first-person psychological state sentences that exhibit the context-shiftiness described above *transparent ascriptions*. Transparent ascriptions seem to tell us something about the speaker's psychological state in some contexts, but in other contexts they primarily tell us something about the world. Philosophers and linguists have offered two general strategies for explaining transparent ascriptions. One strategy, which I will call the *Gricean view*, contends that the meaning of transparent ascriptions never changes relative to context. On this view, psychological verbs invariably contribute information about psychological states to the propositions expressed by transparent ascriptions. But, in appropriate contexts, such psychological state propositions imply hedged information about the non-psychological world, and, in such contexts, this implied information is of central importance. The second strategy suggests that the contribution that certain psychological verbs make to the truth conditions expressed by sentences in which they occur varies. Sometimes these verbs only contribute information about psychological states; sometimes they only, or additionally, serve to hedge the clauses with which they jointly comprise sentences. So, according to this

direct expression view, the meaning of transparent ascriptions changes relative to context.

Here, after briefly explaining how the topic of transparent ascriptions fits into a large debate within linguistics and philosophy, I further explicate the Gricean and direct expression views. Next, I discuss some arguments against each view, and, after showing why each of these arguments fails, suggest that evidence from linguistic processing studies may prove helpful in assessing which view is correct. Finally, I offer three original experiments, which – while far from decisive – provide some initial support for the direct expression view, and, more importantly, suggest that transparent ascriptions present a promising topic for future work in experimental pragmatics.¹

1. The Gricean picture of communication

The debate over transparent ascriptions can be seen as part of a much larger debate over semantic content and the role of pragmatic processing in linguistic understanding. This debate concerns the accuracy of an influential picture of communication that contemporary linguists and philosophers have inherited from twentieth-century philosophers of language – most notably, Paul Grice (1989). In this section, I discuss Grice's picture of communication in detail, and illustrate an important class of argument – 'context-shifting arguments' – that contemporary theorists have used to challenge aspects of Grice's picture. Transparent ascriptions constitute ideal fodder for context shifting arguments.

According to Grice, what is communicated by an utterance in a particular context can be factored into two elements: what is said and what is implicated. What someone says is 'closely related to the conventional meaning of the words (the sentence) he has uttered' (25). What someone implicates follows from rational inferences from what is said, assumptions about the context of utterance, and knowledge of certain conventional maxims of conversation. In addition to maintaining that what is said is what the spoken sentence means "in virtue of the particular meanings of the elements of [the sentence], their order, and their syntactical character" (87), Grice also maintained that what is said is propositional – or, at least, truth-evaluable – in nature. In fact, the propositional nature of what is said is key to Grice's account of implication. The Gricean framework of conversational implicature is one of rational calculation from conventional maxims of communication, which emphasize properties closely tied to truth-conditions, such as accuracy and informativeness. Thus,

on the Gricean picture of communication, what is said is closely related to the conventional meanings of the words that make up the spoken sentence *and* grounds what is implicated in virtue of its propositional nature.² But can these two properties of what is said happily coexist? Over the past several decades, linguists and philosophers of language have vigorously debated this aspect of the Gricean picture. Many theorists have argued against sentence meanings that are both closely related to the conventional meanings of words in the way Grice suggested *and* truth-evaluable, and they have employed so-called context-shifting arguments in making their case.

If what is said corresponds to the words that make up the spoken sentence in the way that Grice maintained, then there can be no change in what is said without a change in the particular meanings of the words that comprise the spoken sentence, their order, or their syntactic character. If what is said is propositional in nature, as Grice also maintained, then, in the case of indicative sentences, what is said determines a set of conditions under which the sentence is true. Thus, if what is said is propositional *and* correlated with the spoken sentence in the way Grice maintained, then there can be no change in the truth conditions of what is said without a change in the particular meanings of the words that comprise the sentence, their order, or their syntactic character (setting aside the small class of traditional indexical and demonstrative expressions, which everyone agrees change their truth-conditional contribution relative to context). According to proponents of context-shifting arguments, the truth conditions of what is said sometimes shift across contexts in which there is no change in the words that comprise the sentence, their order or their syntactic character. If they are correct, then either what is said is not propositional (that is, the truth conditions attributed to what is said actually follow from what is said together with something else) or what is said does not correspond to the words that comprise the spoken sentence, their order and their syntactic character. In either case, the Gricean picture must be modified to allow for context-sensitive processes that work at a sub-propositional level. Either we deny the propositional nature of what is said and allow for processes that work to enrich a non-truth-evaluable sentence meaning (perhaps a proposition schema) into distinct propositions in different contexts. Or else we maintain that what is said is propositional and thus deny sentential correlation by admitting sub-sentential pragmatic processes that modify or supplement the spoken words to generate what is said in different contexts. Either way contextually sensitive pragmatic processes initially transform a non-propositional communicative conveyance into a proposition appropriate to figure into the calculation of implicatures. In

other words, contra Grice, context sensitive pragmatic processes help generate the propositions sentences *directly express*. (These propositions are directly expressed in the sense that they do not follow from and are not logically posteriori to other propositions expressed through the utterance.)

The key premise in context-shifting arguments, which their proponents accept and their opponents deny, is that the truth conditions of what is said sometimes change across contexts in which there is no change in the words that comprise the relevant sentences, their order, or their syntactic character. Various sentences have been offered as context-shifting examples (see Searle 1978 and Travis 1989 for a number of early examples, and Cappelen and Lepore 2005 for a wider discussion), and various theorists concede context-shifting for certain expressions but not others. In effect, the debate has become one over the extent to which the traditional Gricean picture of communication is accurate. More perspicuously, it is a debate over which of various kinds of linguistic expressions are evidence of context shifting and how best to deal with the shiftiness of these expressions. Although they have not been a focus of debate, transparent ascriptions lend themselves to context-shifting arguments, because the truth-evaluable contents these utterances primarily convey seem to change relative to context and such change is not localized to obviously context-sensitive indexical or demonstrative terms.

2. The Gricean view of transparent ascription

Consider a transparent ascription, the sentence “I think Burger Palace is open until midnight”. If Rebecca utters this sentence in a context in which Tina is craving an 11 p.m. snack, there’s a natural sense in which Rebecca would only be taken to speak truthfully if Burger Palace is open until midnight. If, however, Rebecca is explaining the practical syllogism to Tina by describing the mental reasoning that led her to turn right on Northland Avenue at 11 p.m. – if, for example, she is saying, “I want a snack; I think Burger Palace is open until midnight; and I believe Burger Palace is on Northland; thus, I turned right on Northland” – she will be taken to speak truthfully only if she actually *thinks* Burger Palace is open until midnight, regardless of whether it actually is or not. Intuitively, then, what Rebecca has primarily expressed by her utterance of this transparent ascription seems to change across contexts, though the sentence she speaks does not.³ Response patterns to Rebecca’s utterance also favor this conclusion. It would be appropriate to object to Rebecca’s utterance in the first,

snack-seeking context by saying, “No, it’s not. It closes at 10 p.m. on weeknights”. However, this would not be an appropriate rejoinder to the same statement in a psychological description context. It seems, then, that what is primarily expressed by utterances of, “I think Burger Palace is open until midnight”, changes relative to context. And this conclusion generalizes to other transparent ascriptions.

Now, a change in what is *primarily* expressed by utterances of a particular sentence relative to context does not entail a change in what is *directly* expressed by the sentence relative to context. After all, what is primarily expressed by, “Jordan is a bulldog”, can change relative to context – whether we are speaking about my new pet or the bouncer at the local bar, for example. But many would deny that what is directly expressed by this spoken sentence changes across these contexts. Context-relativity in what is primarily expressed by utterances of a particular sentence can be dealt with in either of two ways. It can be identified as context-relativity in what is directly expressed, and thus contribute to context shifting arguments. Or, alternatively, it can be explained as an instance in which what is indirectly expressed by an utterance changes across context.⁴

Adherents to the Gricean view would deny that transparent ascriptions are instances of context-relativity in what is directly expressed by a sentence. They would point out that information about the world can be conversationally implicated when speakers make claims about their psychological states, and that such implications – and not what the pronounced sentences actually say – can, in some conversational contexts (but not others), be of primary importance. For example, Griceans would argue that a claim about the world, *that Burger Palace (may) be open until midnight*, is implicated when a sentence that invariably only expresses a claim about one’s psychological state, “I think Burger Palace is open until midnight”, is uttered in a snack-seeking context. And in such a context, they would maintain, this implication, and not what the spoken sentence actually says, is of primary importance. Since the truth-conditions expressed by transparent ascription sentences do not change on this explanation, this is not an instance of context-shifting.

Kauppinen (2010) offers a succinct statement of this Gricean view.⁵ Of the Gricean account of conversational implicature in general, he writes:

a conversational implicature to the effect that *p* is generated . . . whenever the speaker says (or makes as if to say) that *q* and it is mutually known that the hearer needs to assume that the speaker intended to convey that *p* in order to see the speaker as complying with the Cooperative Principle. (442)

Applying this general account to the present case of transparent ascription, we get the following account of how Tina could work out *that Burger Palace may be open until midnight* from Rebecca's utterance, "I think Burger Palace is open until midnight", in the snack-seeking context (as modified from Kauppinen 2010):

- (1) The mutually acknowledged purpose of the conversation is to acquire information about available snack options.
- (2) Rebecca's utterance would be irrelevant unless the information she conveys about her psychological state were intended to convey some information about available snack options (from 1 and general considerations of relevance).
- (3) By saying what she says, Rebecca only attributes to herself the thought that Burger Palace is open until midnight.
- (4) Assume that Rebecca is abiding by the Maxim of Relation.
- (5) So, in speaking as she does, Rebecca must intend to convey some information about places that are open where one could acquire a snack, specifically, she must mean to convey that what she thinks is true, that Burger Palace is open until midnight (from 2, 3 and 4).
- (6) Given that there is a more direct way for Rebecca to convey information about restaurants that are open, namely a straight assertion like "Burger Palace is open until midnight", she is violating the Maxim of Quantity by being unnecessarily informative in using a psychological verb, unless calling attention to her psychological state plays some role.
- (7) The most obvious role that calling attention to her psychological state plays is to make salient the possibility of divergence between how things appear to one and how they actually are.
- (8) Assume that Rebecca is abiding by the Maxim of Quantity.
- (9) So, in speaking as she does, Rebecca intends to convey that she is not certain that Burger Palace is open until midnight. (from 6, 7 and 8)
- (10) So, in speaking as she does, Rebecca intends to convey information about Burger Palace's hours, namely that the restaurant is open until midnight, but also that she is not certain about that information. (from 5 and 9)

As Kauppinen concludes, "the Maxim of Relation explains why an assertion that is literally about the speaker's psychological states can convey information about the world, and the Maxim of Quantity explains why it is hedged" (443). Importantly, in contexts in which the mutually acknowledged purpose of the conversation differed, this assertion would not convey this information.

According to Kauppinen's Gricean view, the truth conditions of what is said do not change across different contexts in which transparent ascription sentences are uttered. On this view, it is something of a misnomer to think of these as transparent ascriptions at all. Instead, what changes across different contexts is what the propositions expressed by these sentences implicate. Transparent ascriptions pose no threat to the traditional Gricean picture of communication on this explanation, since the Gricean picture maintains that what is implicated and the importance of what is implicated compared to what is said can change relative to context of utterance. I turn now to direct expression views, which, in their more plausible form, entail revision of the traditional Gricean picture of communication.

3. The direct expression view of transparent ascription

At the heart of the direct expression view is the claim that the semantic contribution that psychological state verbs make to the propositions directly expressed by sentences changes in the following way: sometimes these verbs contribute to directly expressed claims about one's psychological state, but other times they (at least additionally) facilitate a directly expressed hedged claim about the world.

The classic version of the direct expression view had its origin in mid-century ordinary language philosophy. For example, in *Philosophical Investigations* Wittgenstein (1958) concludes that "believe" is used in the first-person present to express a hesitant assertion, not "an assertion of hesitancy" (192). And, in an influential paper, J. O. Urmson (1952) writes of:

such philosophical war-horses as *know*, *believe*, and *deduce* . . . that when these verbs are used in the first person of the present tense . . . the assertion proper is contained in the indicative clause with which they are associated, which is implied to be both true and reasonable. (495)

Urmson claims that these 'parenthetical verbs', "function with regard to a statement made rather as 'READ WITH CARE' functions in relation to a subjoined notice, or as the foot stamping and saluting can function in the Army to make clear that one is making an official report" (495–496).

Key to Urmson and Wittgenstein's direct expression view are two ideas, (1) that psychological state verbs makes different semantic contributions in different grammatical constructions, and (2) that in the first-person present

tense they do not contribute to the meaning of the sentence, but rather hedge the force of assertion with which the indicative clause is expressed.⁶ In fact, there are good reasons to doubt this classic version of the direct expression view. As Kauppinen (2010) points out, a syntactically localized version of the view has as a consequence “that I simply cannot attribute a psychological property to myself by using ‘believe’, which is wildly implausible” (441). And as Infantidou (1993) shows, the claim that psychological state verbs used in the first-person present tense do not contribute to truth-conditions does not square with a standard test for distinguishing non-truth-conditional from truth-conditional content. The test asks us to consider conditional statements that contain as their antecedents transparent ascriptions. For example, we are asked to compare:

- (1) If I believe the keys are in the car, Max will put them in the kitchen drawer.

to:

- (2) If the keys are in the car, Max will put them in the kitchen drawer.

The key question is whether the conditions under which the antecedent is satisfied for (1) are the same as for (2). They are not. (1) says Max will put the keys in the kitchen drawer under different conditions than (2). Thus, “I believe the keys are in the car” is not semantically equivalent to “The keys are in the car”, as Urmson and Wittgenstein maintained.⁷

These considerations reveal the implausibility of a direct expression view that claims first-person present indicative uses of psychological verbs invariably modulate the force with which the indicative clause is expressed. But they are consistent with other versions of the direct expression view. Direct expression views claim that psychological state verbs sometimes directly hedge the indicative clauses with which they jointly comprise sentences. There are other plausible accounts that would claim this for transparent ascriptions while denying that the hedge is syntactically triggered. For example, the previous considerations leave open direct expression views that claim that first-person present tense psychological state sentences are sometimes used to directly express *only* hedged claims about the world, so long as such uses are triggered by context of utterance, and do not invariably follow in virtue of syntactic features. For such views could maintain that the non-semantic use is triggered only in specific contexts, thus not invariably activated in the first person, as Kauppinen points out is implausible, or in the particular contexts Infantidou’s test invokes.

A more plausible version of the direct expression view claims that, in certain contexts of utterance, a first-person, present tense psychological state sentence can directly express, and the psychological state verb it contains can serve to hedge, the claim about the world contained in the sentence’s indicative clause. Such a view may or may not also maintain that the utterance simultaneously directly expresses that the utterer is in the psychological state indicated by the main clause. Carston (2002) advocates such a view when she writes of transparent ascriptions that, “the content of the complement clause can, on occasion, be communicated and so qualify as an explicature of the utterance” (211). Relevance theorists, such as Carston, deny a Gricean notion of what is said, that is, they deny a syntactically correlated proposition conveyed by an uttered sentence. Instead, they factor what is communicated by an utterance into explicatures and implicatures, where explicatures are those communicated assumptions that are developments of a logical form encoded by the utterance, and implicatures are communicated assumptions that are not developments of such a logical form. A sentence is seen to convey a logical form, which is insufficient to specify a proposition, but can be systematically fashioned by speakers of a language into a proposition, in virtue of beliefs held common among those speakers, including beliefs about the specific context of communication. Relevance theorists speak of assumptions, not propositions, and this and other terminological differences are not merely terminological, but significant to the theory. Nonetheless, this view maintains that in certain contexts transparent ascriptions can directly express the claim about the world contained in their indicative clauses.

What’s more, relevance theory maintains that any extra processing costs must be justified in terms of positive cognitive effects, or worthwhile differences between one’s pre and post processing representations of the world. So, since Rebecca might have said simply, “Burger Palace is open until midnight”, the extra cognitive costs required to process “I think” must be justified in terms of further, positive cognitive effects. As Infantidou (1993) writes of a similar sentence, “the parenthetical construction *I think* needs to be offset by *extra or different contextual effects*. The different contextual effect in this case is generally a diminished commitment to the proposition expressed” (208). Thus, the inclusion of the psychological verb serves to directly hedge the external world claim contained in the indexical clause on this plausible, relevance theoretic version of the direct expression view.

I do not intend a strong commitment to the relevance theoretic version of the direct expression view. Rather, I take this discussion as an existence proof for a plausible version of the direct expression view. A version of the direct expression

view that maintains that the hedging use of first-person psychological statements is triggered only in certain contexts of use has been defended in the transparent ascriptions literature. If that view is correct, then the propositions directly expressed by first-person psychological state sentences change relative to context, and independently of changes in the meaning, order or syntactic character of the words that comprise the spoken sentences. In that case, the traditional Gricean picture provides an inadequate description of communication when it comes to transparent ascriptions.

Which is correct, the direct expression view or the Gricean view of transparent ascriptions? In the next section, I consider some arguments that bear on this question. But, ultimately, I suggest that the best way to answer this question is through empirical investigation.

4. Arguments concerning transparent ascriptions

Our question comes down to this: Do transparent ascriptions sometimes directly hedge the claims about the external world contained in their indicative clauses, or is such hedged information instead implicated in the relevant contexts? In this section, I review some arguments that bear on this question.

Against the idea that transparent ascriptions directly express information about the external world, Kauppinen (2010) writes that:

Any competent speaker of the language, guided merely by general linguistic knowledge about person and tense transformations, can tell that if 'I believe that *p*', as uttered by N.N., is true at *t*, 'N.N. believed that *p*' will be true at later times. (441)

Here Kauppinen claims that "N.N. believed that *p*" can be inferred from N.N.'s saying, "I believe that *p*", in virtue of general linguistic knowledge and nothing else. It is beyond dispute that competent speakers do – and, therefore, can – infer that "N.N. believed that *p*" from N.N.'s saying, "I believe that *p*". But it is debatable whether they do so by virtue of general linguistic knowledge and nothing else. Competent speakers are typically competent reasoners. And any competent reasoner can infer that if *N.N. believes p* is true at *t*, *N.N. believed p* will be true at later times. Therefore, knowledge that "N.N. believed that *p*" will be true at later times would follow from a hedged assertion by N.N. that *p*, so long as asserting hedgedly entails some degree of belief in the proposition tentatively endorsed – as it would if, as seems likely, the norm of hedged assertion were tentative belief.

Since Kauppinen does not defend the claim that one can tell that "N.N. believed that *p*" will be true at later times in virtue of linguistic knowledge and nothing else, this argument does not challenge direct expression.

Kauppinen also offers a methodological argument against the direct expression view. He quotes Grice's modified Occam's razor, that "senses should not be multiplied beyond necessity". And writes that, "the fewer primitive semantic axioms we postulate, the better we can explain the systematicity and productivity of language" (441). But these objections are misplaced when it comes to the plausible version of the direct expression view, which claims that there is but one primitive semantic axiom for (for example) "think", but that it is either underspecified and capable of different enrichments in different contexts or else does not contribute to semantic content at all in certain contexts.

The previous arguments do not provide compelling reasons to reject the direct expression view. But arguments against the Gricean view of transparent ascriptions are not very compelling either. One might point out that, as Urmson noted, the verb "know" seems to admit of analogous transparent uses as the verbs "believe" and "think". For example, suppose we have searched for the keys to your office for a very long time, both in the car and in the house. Now, we are really getting anxious about finding them. Earlier in our search I said:

(3) I think the keys are in the car.

But now I say:

(4) I just know the keys are in the car.

In such a context, both utterances would primarily express something about the location of the keys, namely, that they are in the car. According to the Gricean view, both utterances do this by directly expressing something about my psychological state, either that I think or else know that the keys are in the car. But now suppose we subsequently find the keys not in the car but between the couch cushions. Well, then I might say:

(5) I just knew the keys were in the car, but I guess I was wrong.

On the Gricean view, I here mysteriously avow that, indeed, I did know the keys were in the car, while at the same time asserting that I was wrong, presumably about the location of the keys. But if knowledge is factive, how can this be?⁸

Against this argument, a Gricean might contend that sentences like (5) operate through protagonist projection. According to this explanation, when I say, "I just knew the keys were in the car", I am projecting into my own past state when I falsely

believed I knew. On this proposal, (5) is like Holton's (1997) sentence, "Sally knew that he would never let her down, but, like all the others, he did". But there is a disanalogy here. Sally would affirm that in the past she really believed that she knew he wouldn't let her down. But one might reasonably deny that he really believed he knew the location of the keys when he uttered (4). Regardless of whether protagonist projection can account for this example, there is independent reason to doubt that I actually attribute knowledge to myself in uttering (4). The utterance sounds more natural when it involves the expression "I just know" rather than "I know". But perhaps "just knowledge" isn't equivalent to true knowledge. Perhaps to say, "I just know", is to say I fall short of knowledge. In that case, expressions involving "just know" do not provide compelling evidence against the Gricean view.

Another potential issue for the Gricean view concerns certain novel utterances that are truth-conditionally equivalent to transparent ascriptions. Remember, according to Griceans, (3) means *that I think the keys are in the car*, and – because that would be irrelevant in a key-searching context – thereby implies *that the keys are in the car*. But were that explanation correct the inferential chain should be preserved in any instance in which (3) is replaced by a semantically equivalent utterance, since conversational implicatures are supposedly non-detachable. But this does not seem to be the case. Consider, for example:

(6) I occupy a truth-conditional, doxastic attitude that the keys are in the car.

Intuitively, (6) would not serve to primarily express a proposition about the external world in relevant contexts (not even on the assumption that all conversational participants understood the utterance). However, Griceans might object that this utterance fails to primarily express the relevant external world claim because it so obviously runs foul of the maxim of manner. (6) is very verbose and obscure. For that reason, conversational participants may stall on the utterance, trying to make sense of why the speaker so obviously flouts the maxim of manner. Thus, they may all together fail to consider the relevance of the utterance to the topic of conversation, and thereby miss the hedged implication.

None of the previous arguments are very compelling. They seemingly leave us at a stalemate between the direct expression and Gricean views. But note that, while the debate between these two views has bearing on a large debate in linguistics and philosophy, both are views of how a specific class of linguistic expressions is processed in relevant contexts. How these expressions are processed is an empirical question. Therefore, it is only natural to recruit experimental methods to help us settle this debate. In the next section, I will try to get the ball

rolling by offering several studies examining people's interpretations of what is said in the case of transparent ascriptions.

5. The experimental pragmatics of transparent ascriptions

Imagine a situation in which Tina is craving a late night snack. She asks Rebecca if any place is still open. Rebecca responds, "*Kate said Burger Palace is open until midnight*". Intuitively, by speaking as she does in this context, Rebecca conveys *that Burger Palace may be open until midnight*. What's more, an inference from what Rebecca says to such a conveyance would be supported by the same mechanisms that the Gricean view of transparent ascriptions invokes; the psychological steps inherent in this inference could be summarized analogously to those enumerated in Section 2 above. The Maxim of Relation could explain how an assertion that is literally about what someone said can be used to make a claim about the world, and the Maxim of Quantity could explain why that claim is hedged.

Since the maxims of communication that the Gricean view appeals to are perfectly general, the same maxims will give rise to the same implication, *that Burger Palace may be open until midnight*, whether Rebecca says, "I think Burger Palace is open until midnight", or, "Kate said Burger Palace is open until midnight". Direct expression views, however, focus on the semantic contribution that psychological state verbs make to the propositions directly expressed by sentences. Since such views are compatible with Gricean implicatures, they allow, "Kate said Burger Palace is open until midnight", to conversationally implicate that *Burger Palace may be open until midnight* in a snack-seeking context. But when, "I think Burger Palace is open until midnight", conveys *Burger Palace may be open until midnight*, it does it in a much different way, according to these views. Thus, a natural way to test these views involves looking at how people treat hedged conversational conveyances in instances involving first-person psychological state implications compared to third-person statement attributions. In three studies, I attempted to draw such comparisons.

5.1 Experiment 1: What is said

5.1.1 Methods

In an initial study, participants (N = 81, male = 57, median age range = 25–34) were assigned to one of two conditions.⁹ In a *first-person state* condition, participants

read four very short conversation descriptions in which a respondent uttered a transparent ascription – a sentence about her or his own psychological state. In a *third-person claim* condition, a different group of participants read about conversations in which a respondent uttered a sentence about what someone else had said. For example, participants in the *first-person state* condition read the following conversation description (among others):

Keys 1st-Person State: Beth is looking for her office keys. She asks Andrew if he knows where they are. Andrew says, “I think they are in the car.”

In the *third-person claim* condition, other participants read this conversation-description:

Keys 3rd-Person Claim: Beth is looking for her office keys. She asks Andrew if he knows where they are. Andrew says, “John told me they are in the car.”

Four psychological state verbs (think, suspect, believe, presume) and four utterance verbs (told, uttered, said, remarked) were used in these descriptions.¹⁰ Changes across conditions for this and all subsequent experiments were minimized and related to changes in these verbs.

All conversation descriptions were designed so that a psychological state sentence or a sentence about what someone else had claimed would be naturally taken to convey information about some other state of the world. What were of interest were participants’ assessments of what each sentence *meant*. In particular, I was interested in whether first-person state sentences might be more readily understood to mean the same thing as hedged assertions. So, directly after reading each vignette, participants were asked about the meaning of what was said. For example, for the above *keys* vignettes, participants were asked:

When Andrew says, [“I think they are in the car”,/“John told me they are in the car”,] does what he says mean, “the keys might be in the car”?

They could respond to this question by choosing yes, no, or not sure. For each participant, I calculated a ‘saying score’ by awarding 1 point each time the participant judged that the relevant hedged, external world claim was part of what the respondent had said, and 0 points each time the participant judged that it was not or that they were not sure. Given that there were four test vignettes, there was a maximal saying score of 4.

In addition to the four test descriptions, four decoy descriptions were also included, which involved speakers uttering external world sentences, such as, “the art museum is on Meade Street next to the library”, and, “we don’t sell

appliances in this store”. I ignored participants’ responses to these descriptions. Descriptions were counterbalanced for order. (All descriptions appear in the Appendix.)

5.1.2 Results and discussion

According to the direct expression view, first-person psychological state ascriptions directly express the hedged external world information that analogous, third-person claim ascriptions merely implicate. According to the Gricean view, both statements implicate the relevant hedged external world information, and neither directly expresses it. Given this difference, if the Gricean view were correct, it would be somewhat surprising to find an asymmetry in assessments of what is said between the two conditions. However, if the direct expression view were correct, such an asymmetry would be expected. In experiment 1, there was a significant difference by utterance type in assessments of whether the relevant hedged claims were part of what was said. Contra the Gricean view, hedged claims were more readily identified with first-person transparent ascriptions than with third-person claim sentences (Figure 9.1).¹¹

In addition to this significant difference in aggregated saying scores, we can examine participants’ responses to individual conversation descriptions. Participants judged that relevant hedged claims about the external world were part of the meaning of what was said much more often for first-person state descriptions than for third-person claim descriptions (Figure 9.2).¹²




Figure 9.1 Mean saying scores




Figure 9.2 Percentage of transparency assessments per vignette

Though far from decisive, these results support the direct expression view of transparent ascriptions over the Gricean view, since they would be less surprising if the direct expression view were correct.¹³ Of course, the high level of assessments that what is said includes hedged claims about the external world in the *third-person claim* condition is *prima facie* incongruent with the direct expression view (as well as the Gricean view). But ordinary speakers' assessments that some proposition is part of what is said do not, in and of themselves, constitute good evidence that it *actually* is part of what is said. Support for the direct expression view follows from the inference that a difference in assessments of what is said would be unexpected if the Gricean view were correct, and not from any claim to the expertise of ordinary speakers in assessing what is said as opposed to what is merely meant.

In response to experiment 1, a Gricean might attempt to explain the asymmetry in participant responses not in terms of a difference in the mechanisms that give rise to the relevant hedged information but rather in terms of a difference in what is actually implicated by the two classes of sentences. Perhaps for some reason or other third-person claim sentences actually fail to convey the relevant external world information that first-person state sentences successfully convey

in my vignettes. In that case, the asymmetry in participants' responses might be indicative of a desire to indicate an asymmetry in what is conveyed, and not a difference in how it is conveyed.¹⁴ With this potential response in mind, I devised a second experiment that assessed whether both kinds of utterances successfully convey the relevant hedged claims.

5.2 Experiment 2: What is said versus what is conveyed

5.2.1 Methods

As in the previous experiment, participants in experiment 2 ($N = 128$, male = 65, median age range = 25–34) were assigned to one of two conditions and presented with four very short conversation descriptions containing either first-person state sentences or third-person claim sentences. As before, participants were asked about the meaning of what was said. However, in this experiment they were first asked whether the relevant hedged external world claim was conveyed by what was said.¹⁵ So, for example, in experiment 2 participants read the following keys description:

Beth is looking for her office keys. She asks Andrew where they are. Andrew says, ["I presume that they are in the car."/"John remarked that they are in the car".]

Directly after reading each description, participants were asked whether hedged information was conveyed by the relevant statement:

When Andrew says, ["I presume that they are in the car."/"John remarked that they are in the car".] does he in some sense convey that the keys might be in the car?

They could respond to this question by choosing yes, no or not sure. Participants were then asked about the meaning of the spoken sentence using the same dichotomous measure as in experiment 1:

When Andrew says, ["I presume that they are in the car."/"John remarked that they are in the car".] does what he says mean, "the keys might be in the car"?

Participants responded to this question by choosing yes, no or not sure.

As before, I calculated a 'saying score' by awarding 1 point each time the participant judged that the relevant hedged, external world claim was part of what the respondent had said, and 0 points each time the participant judged that

it was not or that they were not sure.¹⁶ For each participant, I also calculated a ‘conveying score’ by awarding 1 point each time the participant judged that the relevant hedged, external world claim was in some sense conveyed by what was said, and 0 points each time the participant said that it was not or that they were not sure. Thus, there were maximal saying and conveying scores of 4.¹⁷

Aside from the addition of the conveying probe, careful readers will have noticed several other differences between the keys descriptions in experiments 1 and 2. First, note that whereas Andrew talks about what he *thinks* or what John *told* him in experiment 1, in experiment 2 he talks about what he *presumes* or what John *remarks*. Although the same four psychological state verbs and four utterance verbs were used across both experiments, I altered the descriptions so that different verbs were paired with different conversational topics. This was intended to allay a potential worry that something specific about the particular verb-conversation pairs was responsible for the experiment 1 effect. Next, note that whereas in experiment 1 Beth asks if Andrew *knows* where her keys are, in experiment 2 she simply asks where her keys are. Both the keys and store vignettes from experiment 1 included protagonists who, strictly speaking, could be seen as asking about other protagonists’ mental states. This potentially biasing phrasing was eliminated from the descriptions included in this and subsequent experiments.¹⁸ Finally, the descriptions were also standardized so that each contained an embedded indicative *that*-clause. So, for example, whereas in experiment 1 Andrew says, “I think they are in the car”, in experiment 2 he says, “I presume *that* they are in the car”.

The same four decoy vignettes from experiment 1 were also included in experiment 2, and all vignettes were counterbalanced for order.

5.2.2 Results and discussion

The purpose of experiment 2 was to counter the Gricean response to experiment 1: that third-person claim sentences actually fail to convey the relevant external world information that first-person state sentences successfully convey in my conversation descriptions. While lay-participants may lack mastery of semi-technical terms such as “what is said”, my assumption in experiment 2 was that all ordinary speakers are (roughly) equally adept at understanding what is conveyed by an utterance in context. In support of this assumption is the manifest fact that all ordinary speakers navigate conversations with roughly equivalent aplomb, a skill that philosophers of language agree rests largely on grasping what is conveyed by utterances in context (not merely what is said).

With this assumption in mind, I compared participants’ conveying scores with their saying scores across conditions. My hypothesis was that conveying scores would be high and not significantly different across conditions. I further hypothesized that while saying scores would be – as in study one – high in both conditions, they would significantly differ across conditions, with third-person claim saying scores being lower than first-person state saying scores. These hypotheses were borne out.

I subjected conveying and saying scores to a mixed design 2 (within-subject factor: saying/conveying) $\times$ 2 (between-subject factor: first-person state/third-person claim) ANOVA. I found main effects for both the between factor (first-person state/third-person claim)¹⁹ and the within factor (saying/conveying).²⁰ Importantly, I also found an interaction effect such that people were less likely to judge that relevant hedged claims were part of what was said for *third-person claim* vignettes, than to judge that they were part of what was conveyed for such vignettes, or part of what was said and conveyed in *first-person state* vignettes (Figure 9.3).²¹

Focusing on conveying scores, we can see that participants took both the psychological state sentences and the sentences about what someone else had claimed as they were designed, to convey the relevant hedged information about states of the world. Conveying scores were high and did not significantly differ for *first-person state* and *third-person claim* conditions.²² Saying scores, however, were significantly different across conditions.²³ These findings tell against the Gricean response to experiment 1. It seems that people do think the relevant




Figure 9.3 Mean saying/conveying scores across first- and third-person vignettes

external world information is conveyed both by first-person psychological statements and third-person claim reports in my vignettes, so a difference at this level of communication cannot explain the difference (found in both experiments 1 and 2) between assessments of what is said.

However, though significant, the effect of condition on saying scores was much smaller in experiment 2 than in experiment 1. Additionally, ignoring the aggregate saying scores and looking at each vignette independently, chi-square tests reveal significant differences in participants' responses to the saying-question for only two of the vignettes (keys and store).²⁴ What's more, when those participants who did not think or were not sure that the hedged assertion was conveyed by the relevant utterance in the conversational context are excluded from the analyses (as they arguably should be), a significant difference remains only for keys (though store is approaching significance).²⁵

Why might the saying-effect observed in experiment 2 be smaller than that observed in experiment 1? Perhaps this difference is explained in part by the pull of consistency. What is conveyed, what is said and what is meant are related notions. Perhaps since participants in experiment 2 had to first assess whether relevant hedged information was conveyed by an utterance, and since they overwhelmingly judged that such information *was* conveyed by both kinds of utterance, they felt some pressure to answer the related question concerning the meaning of what was said consistently – to also answer this question in the affirmative. Such an effect could have led saying-assessments to be higher for third-person claim vignettes than they would have otherwise been, and higher than they were in experiment 1, which did not include the conveying probe.

There are several ways to test this consistency explanation. For example, we could alter the order of the elements of the experiment, asking participants the conveying question after they had already answered the saying question. Or we could alter the measure used in the study, choosing a means of response that would not simply emulate the yes/no/not sure response mechanism appearing in study one. In experiment 3, I devised a new scalar measure for the saying question and altered the order of the saying and conveying questions.

5.3 Experiment 3: Convey after say

5.3.1 Methods

As in the previous experiment, participants in experiment 3 ($N = 112$, male = 78, median age range = 25–34) were assigned to one of the two conditions, *first-*

person state or *third-person claim*. Within each condition, participants were presented with the relevant four short conversation descriptions (and four decoy conversation descriptions) used in experiment 2. However, in this experiment, after reading each description, participants were immediately asked about the meaning of what was said using a new saying probe, which elicited level of agreement with a meaning equivalence statement:

How much do you agree with the following statement?

When Andrew says, ["I presume that they are in the car"/"John remarked that they are in the car"] what he says means, "the keys might be in the car".

Participants indicated their level of agreement on a seven-item scale anchored with "Totally disagree" and "Totally agree", and with "Neither agree nor disagree" as the midpoint. Subsequently, after reading each vignette and responding to each saying probe, participants were presented with a new set of instructions informing them that, "Some of the following stories will resemble what you have already read, however, the questions differ", and instructing them to read carefully. Participants then reread the four relevant test descriptions and assessed whether the relevant hedged external world claim was conveyed by what was said. Once again they responded to this question by choosing yes, no or not sure.²⁶

For each participant, I again calculated a 'conveying score' by awarding 1 point each time the participant judged that the relevant hedged, external world claim was in some sense conveyed by what was said, and 0 points each time the participant said that it was not or that they were not sure. Furthermore, I generated an 'overall agreement score' for each participant, by calculating their mean level of agreement with the meaning equivalency statements for the four test vignettes.

5.3.2 Results and discussion

As in experiment 2, participants judged that both first-person state and third-person claim sentences conveyed the relevant hedged information about states of the world. Conveying scores were again high and did not significantly differ for *first-person state* and *third-person claim* conditions.²⁷ However, inline with the previous two experiments, overall scores for the saying probe were significantly different across conditions, with participants agreeing more with the meaning equivalency statements in *first-person state* than in *third-person claim* (Figure 9.4).²⁸




Figure 9.4 Mean overall agreement with meaning equivalence statements

This saying effect size was larger than that observed in experiment 2, and more in line with that observed in experiment 1.

Ignoring overall scores and looking at each vignette independently, we see that participants tended to agree more with the meaning equivalence statements in *first-person state* than in *third-person claim*. T-tests revealed significant differences by condition in participants' responses to the saying-probes for three of these vignettes (keys, store and snack).²⁹ These significant differences remained when those participants who didn't think or weren't sure that the hedged assertion was conveyed by the relevant utterance were excluded from the analyses (Figure 9.5).³⁰

Compared with experiment 2, experiment 3 reveals stronger effects for condition when the saying probe appears before the conveying probe. This suggests that the dampened effect observed in experiment 2 is explained by a desire for consistency on the part of participants.³¹ In any case – and as I discuss in the concluding section below – the findings from the three experiments considered in their totality provide *prima facie* support for the direct expression view of transparent ascriptions over the Gricean view.

6. General discussion and a sketch of future work

The previous experiments reveal a stable difference in participants' understanding of the relationship between, on the one hand, transparent ascriptions (*first-person state* sentences) and hedged information, and, on the other, *third-person*




Figure 9.5 Mean agreement with meaning equivalency statements for those who judged the hedged claim was conveyed

claim sentences and hedged information. Across a variety of measures and conversational contexts, participants associated hedged information with the meaning of transparent ascriptions to a greater extent than with the meaning of *third-person claim* sentences. Despite this stable difference in assessments of what was said, no difference emerged between participants' assessments of what was conveyed. Both *first-person state* and *third-person claim* probes were judged to convey the relevant hedged information at the same high rates. This pattern of responses is entirely unsurprising on the direct expression view, according to which the difference in participants' assessments of what was said is explained by a difference the theory posits in the way in which the two kinds of utterances convey the relevant information. However, the Gricean view posits no such processing difference, so the observed pattern is not explicable from the inner workings of that theory alone. In the absence of convincing deductive arguments, abduction from these experiments favors the direct expression view of transparent ascriptions over the Gricean view, if only weakly. The weakness of this argument to the best explanation is inversely proportional to the plausibility of alternative explanations of the pattern of results. Here, in concluding, I briefly reflect on one possible alternative explanation.

The direct expression view claims that differences in participants' assessments of what is said are reflective of linguistic differences in the way in which two kinds of sentences convey hedged information. It falls to the Gricean to offer a distinct and more plausible account of what the observed asymmetry reflects. One salient possibility attempts to explain the asymmetry in terms not of a linguistic but of an epistemic difference. After all, on the Gricean picture, first-person state sentences report the speaker's belief, whereas third-person claim sentences only report what the speaker heard someone else say. If a speaker is thought to be reliable, his assertion of his own belief that, for example, the keys are in the car will provide reasonably good evidence that the keys might be in the car. However, his assertion that someone else said this is the case, will presumably provide less good evidence regarding the probable location of the keys, since this evidence invokes both the reliability of the speaker and the person he cites. According to the epistemic account of the observed pattern of results, participants' answers to questions about the meaning of what was said reflect (at least in part) their assessments of the quality of information regarding the probable location of the keys. And, since they have lower quality information in the third-person cases, they are less likely to judge that the hedged information is part of what is said in these cases.³²

This alternative, epistemic account of the pattern of saying responses faces difficulties. For one thing, it seems unlikely that the posited epistemic effect would influence saying but not conveying responses. Nonetheless, consideration of individual vignettes may support the epistemic account. My weakest effects were for the bus vignette, in which the interlocutor is just a random bystander, not, as in the other vignettes, a person with whom the inquisitor seems to have a relationship. Perhaps because the informant is unknown to the inquisitor (and is, in turn, conveying information from someone the informant seems not to know), there is perceived to be little difference in the informational value of what is said across conditions for this vignette.³³

Now, it is worth noting that bus saying scores are generally *higher* than saying scores for other vignettes in the *third-person claim* condition, which is not indicative of relatively low informational value assessments. But regardless of the ultimate success of the epistemic account, I welcome alternative explanations for the observed pattern of responses, and I think the argument from individual vignettes is on the right track. Effect sizes across my various experiments seem to follow conversation descriptions, not the particular verbs that appear in the target sentences. This suggests an important role for context in explaining the pattern of responses. Such a role is, of course, perfectly consistent with the plausible direct expression view, which favors a contextual trigger for transparent

uses. But how this trigger works and what exactly it is triggering remains to be spelled out.

Perhaps real headway on the topic of transparent ascriptions will be made only when the more sophisticated methods of experimental pragmatics are brought to bear on it.³⁴ Ultimately, I will be satisfied if I have convinced others that transparent ascriptions are a worthy target of investigation. With luck, experimental philosophers and researchers working within experimental pragmatics will answer my plea and help settle this debate. I believe it is a worthy endeavor.³⁵

Appendix

Conversation Descriptions for Experiment 1:

Keys: Beth is looking for her office keys. She asks Andrew if he knows where they are. Andrew says, ["I think they are in the car"/"John told me they are in the car".]

Snack: Tina is craving a late night snack. She asks Rebecca if any place is still open. Rebecca says, ["I believe Burger Palace is open until midnight"/"Kate said Burger Palace is open until midnight".]

Store: Carl noticed that the prices on shoes at his favorite sporting goods store are very low. He asks Dave if he has any idea why. Dave says, ["I suspect that the store is going out of business"/"Fiona uttered that the store is going out of business".]

Bus: Philip has been waiting five minutes for a bus. He asks the man standing next to him how often the buses run. The man says, ["I presume that they run every ten minutes"/"A passenger remarked that they run every ten minutes".]

Conversation Descriptions for Experiments 2 and 3:

Keys: Beth is looking for her office keys. She asks Andrew where they are. Andrew says, ["I presume that they are in the car"/"John remarked that they are in the car".]

Snack: Tina is craving a late night snack. She asks Rebecca if any place is still open. Rebecca says, ["I suspect that Burger Palace is open until midnight"/"Kate uttered that Burger Palace is open until midnight".]

Store: Carl noticed that the prices on shoes at his favorite sporting goods store are very low. He asks Dave why. Dave says, ["I think that the store is going out of business"/"Fiona told me that the store is going out of business".]

Bus: Philip has been waiting five minutes for a bus. He asks the man standing next to him how often the buses run. The man says, ["I believe that they run every ten minutes"/"A passenger said that they run every ten minutes".]

Decoy Descriptions (Experiments 1, 2, and 3):

Museum: Nate is looking for the art museum. He asks a woman on the street if she knows where it is. She says, "The art museum is on Meade Street next to the library".

Bank: It's Saturday and Martha needs to deposit a check. She asks Evan when the bank closes. Evan says, "The bank closes at noon on Saturday".

Coffee Maker: Heidi is looking for a new coffee maker. She asks a store clerk where the kitchen appliances are. The clerk says, "We don't sell appliances in this store".

Car Lot: Leo is at a used car lot. He asks a man walking by if there's any chance he knows how much the red convertible costs. The man says, "Actually, I'm a salesman here".

Notes

- 1 While my discussion in this chapter will focus on transparent ascriptions involving doxastic state verbs like "believe" and "think", I take it that much of what I write will have wider application to ascriptions involving the large class of psychological verbs considered by Urmson (1952).
- 2 This does not entail that it contributes to the actual psychological process of calculating implicatures (see, e.g. Saul 2002). However, many neo-Griceans and contemporary Gricean critics assume that it typically does (and that the significance of the Gricean account hinges on its psychological reality). I will follow most contemporary participants in this debate by making this assumption.
- 3 A proposition is primarily expressed by an utterance if it is one that hearers would take as a main point of the utterance. I will assume that, if such a proposition were known to be true, hearers would generally judge that the speaker had spoken truly. That a proposition is primarily conveyed by an utterance does not entail that it is directly expressed by the pronounced sentence, since hearers may take an implication of what is directly expressed as a main point of an utterance.
- 4 Or, alternatively, as an instance in which what is of primary importance – what is said or what is implicated – changes across context.
- 5 For a different view in the Gricean spirit, see Jayez and Rossari (2004).

- 6 This classic version of the direct expression view is, in fact, consistent with Gricean tradition.
- 7 Infantidou (1993, 1994) is primarily interested in what she calls 'true parentheticals', sentences in which the psychological state verbs are relegated to a dependent clause, such as, "The keys are, I believe, in the car". In part on the basis of semantic tests such as those discussed above, she concludes that these sentences are syntactically distinct from main-clause parentheticals, like "I believe the keys are in the car". She then offers a detailed account of true parentheticals, according to which they can be used to simultaneously perform two distinct speech acts. Since my concern is with the semantics of main clause parentheticals, I set aside these reflections on true parentheticals, which I agree are syntactically distinct.
- 8 Of course, one might simply deny that knowledge is factive, as some recent philosophers have argued.
- 9 Participants in this and subsequent experiments were recruited and tested using commercially available online platforms (Qualtrics and Amazon Mechanical Turk). All participants were located in the United States. Each was paid \$0.40 for their participation.
- 10 These verbs were selected on the basis of their assessed synonymy with one another in a pre-test.
- 11 $t(78) = 4.98, p < 0.001$ (two-tailed), SE (third-person claim) 0.206, SE (first-person state) 0.064, $r = 0.49$. (One participant was excluded from this analysis due to failure to respond to all vignettes.)
- 12 Keys: $\chi^2(1, N = 81) = 6.32, p = 0.012$; Snack: $\chi^2(1, N = 81) = 12.99, p < 0.001$; Store: $\chi^2(1, N = 81) = 18.8, p < 0.001$; Bus: $\chi^2(1, N = 80) = 4.11, p = 0.043$. (These data collapse "no" and "not sure" responses. For no vignette other than Store $\times$ Third-person claim did more than two participants choose the "not sure" response. For this vignette, nine participants chose "no", and eight chose "not sure".)
- 13 Two comprehension controls were included at the beginning of this experiment. All significant results stand when participants were excluded for failure to pass one or both controls.
- 14 Of course, a Gricean who offered this response would face the burden of explaining why the hedged information is implicated by the first-person state sentence, but not the third-person claim sentence.
- 15 The same four psychological state verbs (think, suspect, believe, presume) and four utterance verbs (told, uttered, said, remarked) used previously were used in these vignettes. However, conversation descriptions were altered so that different verbs were paired with different conversational topics. (This was intended to allay a potential worry that something specific about the particular verb-context pairs was responsible for the previous effect.) The same four decoy vignettes from experiment 1 were also included in this experiment, and all vignettes were counterbalanced for order.

- 16 “Not sure” responses for the saying question: Keys (4), Store (4), Burger (3), Bus (7).
- 17 “Not sure” responses for the conveying question: Keys (0), Store (1), Burger (2), Bus (5).
- 18 Thanks to Wesley Buckwalter for raising this worry in response to an earlier draft of this chapter.
- 19 First-person/Third-person Effect: $F(1, 126) = 6.725$, $p = 0.011$, $\eta^2_{\text{partial}} = 0.051$.
- 20 Saying/Conveying Effect: $F(1, 126) = 11.208$, $p = .001$, $\eta^2_{\text{partial}} = 0.082$.
- 21 $F(1, 126) = 4.076$, $p = 0.046$, $\eta^2_{\text{partial}} = 0.031$.
- 22 Mean conveying scores: first-person state = 3.86; third-person claim = 3.74. $t(126) = 1.29$, $p = 0.2$ (two-tailed), SE (first-person state) 0.055, SE (third-person claim) 0.074, $r = 0.11$.
- 23 Mean saying scores: first-person state = 3.76; third-person claim = 3.35. $t(126) = 2.727$, $p = 0.007$ (two-tailed), SE (first-person state) 0.067, SE (third-person claim) 0.132, $r = 0.24$.
- 24 Keys: $\chi^2(1, N = 128) = 6.79$, $p = 0.009$; Snack: $\chi^2(1, N = 128) = 2.32$, $p = 0.13$; Store: $\chi^2(1, N = 128) = 5.18$, $p = 0.023$; Bus: $\chi^2(1, N = 128) = 0.67$, $p = 0.41$. (As above, this and the next analysis collapse “no” and “not sure” responses. See note 16 for this data.)
- 25 Keys: $\chi^2(1, N = 125) = 7.52$, $p = 0.006$; Snack: $\chi^2(1, N = 122) = 2.14$, $p = 0.14$; Store: $\chi^2(1, N = 121) = 3.44$, $p = 0.064$; Bus: $\chi^2(1, N = 118) = 0.16$, $p = 0.689$.
- 26 I used different saying and conveying measures in this experiment to minimize the possibility of a reverse consistency effect, whereby assessments of the now earlier-occurring saying probe might influence responses to the subsequent conveying probe.
- 27 Mean conveying scores: first-person state = 3.71; third-person claim = 3.73. $t(108) = 0.128$, $p = 0.899$ (two-tailed), SE (first-person state) 0.085, SE (third-person claim) 0.144, $r = 0.013$. (Two participants were excluded from this analysis for failing to answer the conveying probe for all vignettes.)
- 28 Mean overall agreement scores: first-person state = 5.89; third-person claim = 5.06. $t(110) = 3.99$, $p < 0.001$ (two-tailed), SE (first-person state) 0.097, SE (third-person claim) 0.188, $r = 0.35$.
- 29 Keys: $t(110) = 2.69$, $p = 0.008$ (two-tailed), SE (first-person state) 0.129, SE (third-person claim) 0.242, $r = 0.25$; Snack: $t(110) = 4.684$, $p < 0.001$ (two-tailed), SE (first-person state) 0.134, SE (third-person claim) 0.276, $r = 0.4$; Store: $t(110) = 3.271$, $p = 0.001$ (two-tailed), SE (first-person state) 0.153, SE (third-person claim) 0.241, $r = 0.29$; Bus: $t(110) = 1.068$, $p = 0.288$ (two-tailed), SE (first-person state) 0.151, SE (third-person claim) 0.200, $r = 0.1$.
- 30 Keys: $t(106) = 2.4$, $p = 0.018$ (two-tailed), SE (first-person state) 0.128, SE (third-person claim) 0.229, $r = 0.22$; Snack: $t(100) = 4.216$, $p < 0.001$ (two-tailed), SE (first-person state) 0.129, SE (third-person claim) 0.268, $r = 0.38$; Store: $t(102) = 3.203$,

$p = 0.002$ (two-tailed), SE (first-person state) 0.135, SE (third-person claim) 0.22, $r = 0.3$; Bus: $t(99) = 0.199$, $p = 0.842$ (two-tailed), SE (first-person state) 0.164, SE (third-person claim) 0.164, $r = 0.02$.

- 31 Here it is worth pointing out that two unreported experiments, which had the same conveying-saying probe order as experiment 2, but used meaning equivalence saying probes like those appearing in experiment 3, returned the same general pattern of results reported above. For both of these experiments, however, observed saying effects were smaller, and mean responses were higher overall. This suggests that, while the consistency effect observed in experiment 2 can be dampened by varying the type of probe used, it is best dealt with by changing the order of the probes.
- 32 I am thankful to Daniel Cohnitz for suggesting this interesting alternative explanation in a referee report on an earlier version of this chapter.
- 33 I am thankful to Daniel Cohnitz for suggesting this argument.
- 34 For example, experimental methods such as processing time studies, eye-tracking studies and developmental studies have proven useful in unraveling related questions concerning so-called scalar implicatures. See Phelan (2014) for discussion.
- 35 I'm grateful to Wesley Buckwalter, Daniel Cohnitz, Jussi Haukioja, Shen-yi Liao and Thomas Ryckman for reading an earlier version of this chapter and offering many important suggestions and searching criticisms. Their comments improved this work immensely.

References

- Austin, J. L. (1962), *How to Do Things with Words*, 2nd edition. Ed. J. O. Urmson and M. Sbisá. Cambridge, MA: Harvard University Press.
- Bach, K. (2002), ‘Seemingly Semantic Intuitions’, in J. Keim Campbell, M. O'Rourke and D. Shier (eds), *Meaning and Truth*. New York: Seven Bridges Press, pp. 21–33.
- Carston, R. (2002), *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell.
- Grice, P. (1989), *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.
- Holton, R. (1997), ‘Some telling examples: A reply to Tsohatzidis’, *Journal of Pragmatics*, 28, 625–628.
- Infantidou, E. (1993), ‘Parentheticals and relevance’, *UCL Working Papers in Linguistics*, 5, 193–210.
- Infantidou, E. (1994), *Evidentials and Relevance*. PhD thesis, University of London.
- Jayez, J. and Rossari, C. (2004), ‘Parentheticals as conventional implicatures’, in F. Corblin and H. de Swart (eds), *Handbook of French Semantics*. Stanford: CSLI, pp. 211–229.

- Kauppinen, A. (2010), 'The pragmatics of transparent belief reports', *Analysis*, 70 (3): 438–446.
- Phelan, M. (2014), 'experimental pragmatics: An introduction for philosophers', *Philosophy Compass*, 9 (1), 66–79.
- Saul, J. (2002), 'What is said and psychological reality: Grice's project and relevance theorists' criticisms', *Linguistics and Philosophy*, 25, 347–372.
- Urmson, J. O. (1952), 'Parenthetical verbs', *Mind*, 61, 480–496.
- Wittgenstein, L. (1958), *Philosophical Investigations*. Oxford: Blackwell.

Index

- ambiguity 142–151, 166–169
- Beebe, J. 69
- belief ascriptions 49–53, 173–197
- "Benjaminus" case 149–151
- Braisby, N. 121
- Buckwalter, W. 50
- Cappelen, H. 28n, 56n, 90
- Carston, R. 181
- "Catoblepas" case 146–149, 160–161
- causal theory of reference 17, 33, 69–73, 88–90, 159–169
- Chomsky, N. 36, 79
- "Cicero" case 13
- Cohnitz, D. 1, 4, 92–94, 97
- "Columbus" case 11, 20, 45
- communication 91–94, 174–176
- context-shifting arguments 176–179
- conversational implicature 152–153, 177–179
- corpus, the 46–49
- Davidson, D. 67–68, 78
- DeBlanc, M. 133n, 158
- descriptivism 11–14, 33, 69–73, 88–90, 117–120, 140–142, 159–166
- Deutsch, M. 3, 90–91, 114
- Devitt, K. 50
- Devitt, M. 3, 19–20, 24, 27n, 33–46, 78, 90, 93, 97–99, 115–116, 168–169
- dispositions, linguistic 92–104
- Donnellan, K. 44
- Dummett, M. 35
- "Einstein" case 11, 20, 45
- elicited production 49–51
- essentialism, psychological 120–126
- Evans, G. 163–166
- Expertise Defense *see* intuitions, of experts
- extension 66–73, 75–79
- externalism 74–75, 91–92, 120–126
- eye-tracking 100–104
- "Feynman" case 13
- Field, H. 68
- Franks, B. 121
- Genone, J. 47–49, 58n, 125–126, 159–160
- Gernsbacher, M. 101
- "Gödel/Schmidt" case 7–20, 45, 69–73, 89–91, 95–97, 102–103
- Grice, P. 174–179
- Häggqvist, S. 4–5
- Hampton, J. 121
- Haukioja, J. 92–94, 97, 122–125
- Hintikka, J. 34
- Horwich, P. 68
- hybrid theory of reference 33, 123–125, 162–166
- "Ibn Kahn" case 165
- Ichikawa, J. 18–19, 26n, 133n
- implicit scare quotes, problem of 51–53
- Infantidou, E. 180–181
- internalism 74–75, 91–92, 120–126
- interpretation 92–99
- intuitions,
 as a priori 35–36
 as evidence in philosophy 8–11, 17–25, 34–39, 92–99
 of experts 40–42, 66–67, 115–116
 fanciful vs. humdrum 19, 44
 linguistic vs. metalinguistic 97–99
 memory judgments as 43
 perceptual judgments as 39–42
 in thought experiments 43–45, 95–99
 as the 'voice of competence' 36–37
- "Jonah" case 69, 141–142, 144
- Jylkkä, J. 122–125