

Relative truth and the first person

Friederike Moltmann

Published online: 15 April 2009
© Springer Science+Business Media B.V. 2009

Abstract In recent work on context-dependency, it has been argued that certain types of sentences give rise to a notion of relative truth. In particular, sentences containing predicates of personal taste and moral or aesthetic evaluation as well as epistemic modals are held to express a proposition (relative to a context of use) which is true or false not only relative to a world of evaluation, but other parameters as well, such as standards of taste or knowledge or an agent. I will argue that the sentences that apparently give rise to relative truth should be understood by relating them in a certain way to the first person. More precisely, such sentences express what I will call ‘first-person-based genericity’, a form of generalization that is based on an essential first-person application of the predicate. The account differs from standard relative truth account in crucial respects: it is not the truth of the proposition expressed that is relative to the first person; the proposition expressed by a sentence with a predicate of taste rather has absolute truth conditions. Instead it is the propositional content itself that requires a first-personal cognitive access whenever it is entertained. This account, I will argue, avoids a range of problems that standard relative truth theories of the sentences in question face and explains a number of further peculiarities that such sentences display.

Keywords Relative truth · First person · De se · Predicates of taste · Genericity · Propositional attitudes

In recent work on context-dependency, it has been argued that certain types of sentences give rise to a notion of relative truth. In particular, sentences containing predicates of personal taste and moral or aesthetic evaluation as well as epistemic

F. Moltmann (✉)
Directrice de recherche, Institut d’Histoire et de Philosophie des Sciences, et des Techniques
(IHPST), 13 rue du Four, 75006 Paris, France
e-mail: fmoltmann@univ-paris1.fr

modals are held to express a proposition (relative to a context of use) which is true or false not only relative to a world of evaluation, but other parameters as well, such as standards of taste or knowledge or an agent. Thus, a sentence like *chocolate tastes good* would express a proposition p that is true or false not only at a world of evaluation, but relative to the additional parameter as well, such as a parameter of taste or an agent.

I will argue that the sentences that apparently give rise to relative truth should be understood by relating them in a certain way to the first person. More precisely, such sentences express what I will call *first-person-based genericity*, a form of generalization by which the speaker quantifies over every one x in the relevant domain as someone he identifies with, allowing the predicate to apply to x as if it applied to the speaker himself. This account differs from standard relative truth theories in crucial respects: it is not the truth of the proposition expressed that is relative to the first person; the proposition expressed by a sentence with a predicate of taste rather has absolute truth conditions. Instead it is the propositional content itself that requires a first-personal cognitive access whenever it is entertained. Thus, if two agents disagree about a sentence like *chocolate tastes good* without either ‘being at fault’, this is disagreement about a sentence that has absolute truth conditions and thus is either true or false. But the content of the sentence can be grasped (and thus entertained and evaluated) only in an essential first-personal way, namely by applying the predicate to everyone in the domain as if to oneself (and thus allowing two agents to have different first-personal grounds for applying the predicate). On the proposed account, it is conditions on grasping the content, rather than truth conditions what is peculiar to sentence giving rise to intuitions of relative truth.

This account, I will argue, avoids a range of problems that standard relative truth theories face and explains a number of further peculiarities that such sentences display. The account goes along with an independently motivated view on which truth conditions of sentences are not given by mind-independent propositional contents, but rather by mind-dependent ‘attitudinal objects’ (as I will call them), or kinds of such objects, that is, objects of the sort John’s belief that S or the belief that S .

First-person-based genericity is expressed explicitly in English by sentences that contain the generic pronoun *one* as in (1a) or its empty counterpart, so-called arbitrary PRO (as generative syntacticians call it), as in (1b):

- (1) a. One can see the picture from the entrance.
 b. It is possible PRO_{arb} to see the picture from the entrance.

Sentences (1a) and (1b) have a natural reading on which they express a generalization on the basis of the speaker’s own, perhaps unique, experience or action. That is, (1a, b) are naturally used as an expression of the speaker’s own ability to see the picture from the entrance and at the same time express a generalization: for every normal x , x can see the picture from the entrance.

The key observation is that the same intuitions of relative truth that are displayed by sentences with evaluative predicates or epistemic modals arise with sentences such as (1a) and (1b). In previous work I have developed an account of generic *one*

on which generic *one* expresses (contextually restricted) quantification over individuals insofar as the relevant agent identifies with them. In this paper, I will argue that this analysis can and should be carried over to sentences with predicates of personal taste, as well as possibly other sentences containing predicates of aesthetic and moral evaluation or epistemic modals. I take first-person-based genericity to be the source of faultless disagreement in general, for any expressions with which it may arise.

However, not all types of sentences for which a relativist treatment has been proposed are to be analysed in terms of first-person-based genericity. Expressions that involve first-person-based genericity need to be distinguished from expressions that are to be treated in just the same way as *de se* interpreted pronouns (such as adjectives like *left* or *local*). The latter do not give rise to faultless disagreement and do not involve first-person-based genericity.

I will start by presenting two core intuitions that motivate a relativist account, and point out further crucial features of the relevant linguistic expressions. I then give a general account of *de se*-interpreted pronouns and related expressions. After that I introduce the notion of first-person-based genericity on the basis of my previous analysis of generic *one* showing that generic-*one* sentence exhibit the intuitions of relative truth in the same way. I finally argue that the same semantic analysis can and should be carried over to other sentences displaying intuitions of relative truth.

1 The intuitions of relative truth

1.1 Faultless disagreement

One main motivation for a relativist account has been the possibility of faultless disagreement (Koelbel 2002, 2003). Faultless disagreement consists in a situation in which two agents disagree about the truth of a proposition, with neither apparently being at fault. Thus, below, neither A nor B may be at fault, yet they disagree:

- (2) A: Frog legs taste good.
B: No, I disagree, frog legs do not taste good.

In fact A and B subsequently may enter into a dispute, let's say whether they should start putting frog legs on the menu of their restaurant. Faultless disagreement arises not only in a situation of conversation, but also manifests itself in intuitions about two agents being involved in different conversations (MacFarlane 2007) or about two agents' beliefs (Koelbel 2003). What is crucial about faultless disagreement is that both agents seem to be right in their claims or beliefs, but yet they disagree. What is important is that 'to be right', or 'not to be at fault' means more than having a reason or epistemic grounds: the intuition is that both the statement and its negation are in some sense true, though as uttered by different agents.

Faultless disagreement does not arise with sentences that express different propositions when uttered by different speakers. Moreover, faultless disagreement does not arise when the 'judge' is made explicit, as in (3):

- (3) A: Frog legs taste good to me.
 B: Frog legs do not taste good to me.

Related to the possibility of faultless disagreement is the observation that two agents may *agree* about the content of an evaluative sentence even if it is clear that the relevant parameters of evaluation of the two agents are different. Suppose John is a great wine connoisseur, with a comparative experience of various types of wine, whereas Mary has not tasted any wine before, then the parameters evaluation for the proposition that the wine tastes good will certainly be different for John and Mary, yet they may agree that the proposition is true:

- (4) John and Mary agreed that the wine tastes good (but for very different reasons).

This intuition is suitably called *faulty agreement*. Like faultless disagreement, it arises in just the same way in relation to different agents' beliefs. That is, (4) would be true even if John and Mary did not talk to each other, but just believe, for different reasons, that the wine tastes good.

For sentences with predicates of taste, the object of disagreement or agreement is a propositional content that is truth-conditionally incomplete in that it does not involve the relevant judge. Or so the data have been interpreted by relativist theories.

There is another property that has been related to relative truth, namely retraction: the possibility of withdrawing a propositional content once correctly held in the past. I will put this property aside for the main part of the paper and turn to it only briefly at the end.

1.2 Sharing

By 'sharing' I mean the possibility of sharing of propositional contents by agents involved in different contexts of evaluation. Sharing is a property not generally considered in the recent literature on relative truth, and thus the data in this section are new. Sharing has been already discussed at length, though, in regard to moral predicates by Schiffer (1990). If a sentence S involves relative truth, then sharing of propositional contents consists in the intuition that agents, even if they are clearly involved in different contexts of evaluation, share the same propositional content when they have a propositional attitude that S. There are various linguistic manifestations of sharing.

One of them is the possibility of sentences with a conjunctive or plural subject and a single clausal complement of an attitude verb¹:

- (5) John and Mary/These two people believe that the wine tastes good.

¹ The use of *believe* in (5a) suggests that the agent does not himself have the experience. To get the non-subjective first-person-based reading, *know* would be the right verb:

(1) John knows that white chocolate tastes good.

But then faultless disagreement is actually hard to get, see Sect. 5.

(5) requires a single interpretation of the clausal complement, which can thus not involve two different taste parameters or agents.

Another linguistic manifestation of sharing is the validity of the inference in (6), assuming that A's and B's criteria for evaluating wine are known to be quite different (A, but not B, lets say, being a connoisseur):

- (6) A believes the wine tastes good.
B believes the wine tastes good.
 A and B believe the same thing.

Such an inference holds with any propositional attitude or speech act verb. Evaluative predicates, in licensing the inference, thus differ from other context-dependent expressions such as demonstratives, with which the inference is not valid:

- (7) John believes that Mary is there (in New York).
Bill believes that Mary is there (in Boston)
 John and Bill believe the same thing.

The validity of the inference in (7) does not hinge on some 'looseness' of uses of the expression *the same thing*. With a conclusion containing a free relative as in (8) the same sort of inference is validated, assuming again that John's and Mary's taste parameters are rather different:

- (8) John believes what Mary believes, namely that the wine tastes good.

Even more importantly, the same sort of inference is valid with a conclusion containing the corresponding nominalization of the verb, such as *the same belief* or *the same claim*:

- (9) a. John believes that the wine tastes good
Mary believes that the wine tastes good.
 John and Mary have the same belief (share a belief).
 b. John claimed that the wine tastes good
Mary claimed that the wine tastes good.
 John and Mary made the same claim.

These inferences are valid even if John's and Mary's taste parameters' are different (and even are known to be different).

Such 'linguistic' criteria for sharing are not the only criteria for establishing shared meaning when different contexts of assessment are involved. There are also a range of other intuitions that indicate that the meaning, that is, 'conceptual role' of evaluative predicates when used by different agents is the same even if the agents' criteria for applying the predicates are rather different. This point has been made rather thoroughly, in the context of a discussion of moral predicates, by Schiffer (1990). Schiffer discusses various intuitions that show that moral predicates have the same meaning for two agents even if the agents differ in their moral criteria and principles for applying the predicates. For example, even if two agents have different moral principles, their concepts of good and wrong are the same given the roles of those concepts in determining how the agents want the world to be, given the kinds of emotional responses they trigger, given the relation of those concepts to

related ones (*ought, just*), given that the concepts involve the same process of moral training (of punishment and reward), and given the predicates' common 'point', in getting people to behave in a certain way. Needless to say, similar criteria for the individuation of conceptual roles can be found for other evaluative predicates as well.

The crucial point that faultless disagreement and sharing appear to establish is that the propositional content of sentences with predicates of personal taste is the same even when the context-dependent criteria of evaluation involved (such as standards of taste etc.) are clearly distinct.

1.3 Other kinds of expressions exhibiting the intuitions of relative truth

Epistemic modals are another often discussed class of expressions displaying the intuitions of relative truth (Egan et al. 2005; Egan 2007; MacFarlane 2008). They clearly display faultless disagreement: one person may believe or claim that John may be in Paris, while another person with more knowledge disagrees, with neither being at fault. Epistemic modals also lead to sharing of content, as shown by the validity of the following inference:

- (10) Mary believes that it may rain (because she heard the weather forecast).
John believes that it may rain (because he noticed the cloud formation).
 John and Mary believe the same thing (but for different reasons).

The inference is also valid with the conclusion in (11a) or, as faulty agreement, with the conclusion in (11b):

- (11) a. John and Mary have the same belief.
 b. John and Mary agree that it may rain.

There are other context-dependent expressions which give rise to sharing, but not faultless disagreement (or faulty agreement). These are certain temporal and spatial expressions such as *right, left, local or neighboring*.² Thus, from (12a) (which itself involves sharing already), one can infer both (12b) and (12b') even if what counts left for Mary is right for John:

- (12) a. John and Mary both believe that the tree is to the left.
 b. John and Mary believe the same thing.
 b'. John and Mary share the same belief.

² It appears that also certain presupposition triggers allow for sharing, for example expressions like *another time, a second time, or again*:

- (1) John believes that he won in 2006.
Mary believes that she won in 2005.
 John and Mary hope to win another time/a second time/again.
 (2) a. John and Mary hope for the same thing.
 b. John and Mary have the same hope.

Such presuppositional expressions are generally held to be anaphoric to some element in a preceding textual or mental representation. But obviously, the identity of the antecedent element does not bear on the identity of the content. The present paper will not try to offer an account of the phenomenon.

In that situation, however (that is, if what counts left for Mary is right for John), it would not be the case that John and Mary agree. Rather, they clearly disagree. Thus, *left* does not give rise to faulty agreement. Relational adjectives like *left* or *neighbouring* in fact pattern like pronouns interpreted *de se*, and as such they need to be sharply distinguished from predicates of taste and other expressions of the kind. *De se* interpreted pronouns exhibit sharing, as seen from the validity of the inference in (13a), with the two occurrences of *he* being interpreted *de se*, as well as the same inference with (13b) as conclusion:

- (13) a. John thinks that he is the winner.
 Bill thinks he is the winner.
 John and Bill think the same thing.
 b. John and Bill share the same thought.

I will distinguish expressions giving rise to faultless disagreement as well as sharing as *type 1 expressions* from those that give rise only to sharing, which I will call *type 2 expressions*³.

Later we will see that two further features distinguish the two kinds of expressions: type 1 expression exhibit what I will call obligatory and a quasi-first-person-orientedness, whereas type 2 expressions exhibit optional sharing and no quasi-first-person orientedness. We can thus summarize the features of the two kinds of expressions as follows:

- (14) type 1 expressions (predicates of taste, epistemic modals, predicates of aesthetic evaluation, generic *one*):
 faultless disagreement, (obligatory) sharing, quasi-first-person orientedness.
 type 2 expressions (*de se* interpreted pronouns, relational adjectives (*right*, *left*, *local*, *neighbouring*):
 (optional) sharing, no faultless disagreement, no quasi-first-person-orientedness.

2 Standard relativist theories and their problems

The relativist account takes the content of sentences giving rise to relative truth intuitions to be truth-conditionally incomplete. In order to account for faultless disagreement in particular, it relativizes the truth value of such a content to the context of the assessor, in particular the assessor himself or other parameters that belong to the context of the assessor (such as a parameter of taste). Truth then is relative in the following sense:

³ Egan (2007) discusses the same distinction with a sentence like *my pants are on fire* and sentences with epistemic modals. He argues that *my pants are on fire* cannot be asserted when expressing a self-locating proposition and gives a pragmatic account of why this is so. It is fairly clear, however, that the distinction is not a pragmatic one but resides strictly in the formal distinction between type 1 and type 2 expressions (expressions like pronouns interpreted *de se* as opposed to expressions giving rise to relative truth).

(15) Relative Truth

The truth of a sentence *S* is relative iff for an utterance context *u*, the proposition that *S* expresses in *u* is true or false only relative to a context *c* that contains not only a world (and time) but other elements as well.

Here *c* is simply the context to which the truth value of the sentence is relativized (also often called an ‘index’, to distinguish it from the utterance context). McFarlane distinguishes between the context of evaluation, the context that can be shifted by modal operators, and a context of assessment, which is the context of the person evaluating the sentence as true or false, and which thus would correspond to the context *c* in (15). Other authors do not make the distinction between context of evaluation and context of assessment, but take sentences with predicates of taste to have a truth value relative to a context of evaluation containing also a judge (thus extending the ‘index’) (Lasersohn 2005). For the discussion to follow, the difference will not be crucial.

There are different views as to what the context may contain to which the truth of the sentence is to be relativized. On one view, the context may contain whatever parameters seem necessary for the evaluation of the expression in question, such as parameters of taste or epistemic standards (MacFarlane 2005a, b). On another view, it will just contain, besides a world (and perhaps time) of evaluation, the relevant agent, or equivalently, a centered world (a pair consisting of a world and an agent) (Egan et al. 2005; Lasersohn 2005). It appears that a relativist account, especially on the first version, is problematic in several respects.

2.1 Explaining disagreement

The most important problem for the relativist account is that it does not really explain faultless disagreement. Competent speakers will know if the truth value of a sentence they utter is to be evaluated relative to a particular kind of context of assessment. They will know that they mean the utterance of such a sentence to be true relative to their own context. The problem then is, why on the relativist account can there be disagreement among two speakers when the speakers know that the content of their utterances can be both true, though relative to different contexts? If the truth conditions of the sentence are clearly different relative to the two speakers, then this should correspond to a difference in subject matter, rather than to a single content about which there could be disagreement.

Also, not every version of the relativist account seems to be able to explain the difference between *frog legs taste good to me* on the one hand and *frog legs taste good* on the other hand. On one proposal (Lasersohn 2005), the context of evaluation depends in fact on the speaker’s own choice, his taking a particular perspective for evaluating the proposition. If the context of assessment depends on an agent’s choosing a perspective of evaluation, then one would think that this choice of a perspective is part of an overall contextually completed intended meaning of the utterance. But then such a contextually completed meaning would not be relevantly different from the meaning of *frog legs taste good to me* (or *frog*

legs taste good to d, where *d* is the relevant judge). But we have seen that the latter does not give rise to faultless disagreement, whereas the former does.⁴

Another problem for the account concerns the parameters themselves that are used for the relativization of truth. Those parameters, be they agents or evaluative or epistemic standards, are individuated externally and should thus be equally accessible and identifiable for all the interlocutors, given the circumstances of the utterance situation or propositional attitude. Given the standard view about assertion and belief as aiming for truth, why shouldn't a truth-conditionally complete content including the relevant parameters be identified in a given context by each one of the interlocutors and understood as the content to be communicated? That is, why should agreement or disagreement not pertain to such a truth-conditionally completed content as it would serve the aim of assertion or belief, rather than the truth-conditionally incomplete content expressed by a sentence relative to a context of utterance?^{5,6}

The relativist account is also difficult to maintain for the content of a sentence like *frog legs taste good* when it acts as the object of belief. There are two options of how to construe the object of belief in this case. First, the object of belief could be viewed as a pair consisting of a proposition and a context of evaluation (Lasersohn 2005). But this does not seem to do justice to the intuition that two people, one of whom believes that frog legs taste good and the other of whom believes that frog legs do not taste good, disagree rather than having different but compatible overall beliefs. Alternatively, the content of belief could be taken to be a truth-relative proposition with the judge being identified with whoever is the believer (Egan et al. 2005). But again this makes it difficult to account for how disagreement could arise.

To summarize, the relativist account straightforwardly explains sharing. But it does not provide a real explanation of why the relevant examples give rise to disagreement.

2.2 The meaning-intention problem

The first version of the relativist account posits parameters, such as parameters of taste, as part of the context of assessment. The problem with such entities is the same as in a number of cases in which entities of a certain sort are posited as implicit arguments, namely what Schiffer (1987) has discussed as the meaning-intention problem. A speaker generally would not be able to identify entities such as parameters of taste as part of a context relative to which he intends his utterance to

⁴ On MacFarlane's (2005a, b) account, the context of assessment is uniquely determined by whoever assesses the truth value of the proposition. This obviously would not predict an identity of the intended meaning of *white chocolate tastes good* and *white chocolate tastes good to d* (for *d* being the relevant judge).

⁵ Stephenson (2007) attributes the disagreement in faultless disagreement to the aim of (some) assertions to establish a sort of single shared judge as part of the common ground in the context of conversation. The problem of such a discourse-related explanation of faultless disagreement is that it could not be carried over to faultless disagreement that arises with beliefs.

⁶ MacFarlane (2003) actually takes the intuitions of relative truth to require a modification of the notion of assertion, truth not being the aim of assertion anymore.

be made. A speaker may be entirely justified in uttering a sentence like *this tastes good* without in any way being able to identify a ‘taste parameter’ as part of his intentions. This problem of cognitive accessibility of the parameters of a context of assessment is a natural generalization of Schiffer’s meaning-intention problem.

2.3 The *de se* status of the parameters of assessment

Relativist accounts allow for readings involving error through misidentification that in fact do not exist. Suppose that Joe reads a description of someone that he fails to realize is a description of himself and that misdescribes him as someone that likes frog legs when Joe in fact does not. In this case, given a relativist account, *Joe believes that frog legs taste good* would be true (though Joe’s belief would not). An agent *d* could not believe *p* to be true relative to a taste parameter that *d* does not believe to be his own, that is, *d* could not believe *p* to be true relative to *d* without recognizing that *d* is himself. In other words, an agent must identify himself as the judge when standing in an attitudinal relation to an evaluative propositional content. The problem for the relativist account is that it assigns to the additional parameter of assessment a status of a value of a pronoun interpreted *de re*, when in fact the parameter should have the status of the semantic value of a pronoun interpreted *de se*.⁷

2.4 Predictions about attitude contexts

The standard relativist account faces potential problems for attitude reports. McFarlane (2005b) assumes that the additional parameter of assessment of the embedded proposition always becomes the parameter of the assessment of the entire sentence. This is adequate, however, only for non-attitudinal embedding predicates, such as modal and temporal ones, as in (16):

(16) It could be the case that frog legs taste good.

In (16), the taste parameter is clearly that of the speaker. This is different with attitude reports:

- (17) a. John thinks that frog legs taste good.
b. John thinks that it might rain.

Clearly the standard of taste in (17a) and the epistemic state in (17b) are those of John, not the speaker. In general, with embedding attitude verbs, the ‘additional parameter’ will be the one of the described agent. This is not a problem for a relativist account as such, however. As Egan et al. (2005) make clear, for a truth-relative proposition *p*, an attitude report *d believes that S* can be made to be understood as: *d believes S to be true relative to d*.⁸

⁷ For observations and a proposal concerning *de se* status in the case of epistemic modals see Stephenson (2007), fn 15.

⁸ See Stephenson (2007) for an equivalent account.

Lasersohn's (2005) account also gets attitude reports right. Lasersohn takes attitude verbs to express three-place relations between an agent, a sentence content, and a context. The context is the perspective the agent adopts when taking the attitude toward the sentence content and in general it is the context of the agent. But it need not be: the agent may take a different stance, which is the quasi-first-person orientation which I will turn to next.

2.5 The quasi-first-person orientation

There is another important feature of evaluative predicates and epistemic modals which is that their first-person-orientation in independent contexts need not be strict, but may relate to another agent with whom the speaker only identifies. For example, a mother may persuade a child to eat by uttering (18), without thereby expressing her own taste judgment:

(18) Apple sauce tastes good.

The same point can be made with *because*-clauses and questions⁹:

(19) a. John took another spoonful because it tasted so good.
b. Does this taste good?

Here the speaker may just identify with (or project himself onto) John (in 19a) or the addressee (in 19b), without being interested in his own taste judgments at all.

Also, with epistemic modals the first-person orientation may involve an identification with another agent. Thus Egan et al. (2005) note that (20) could be uttered by a speaker who knows better but identifies himself with a person trying to find a way out of a maze:

(20) The exit may be this way.

Let me call the first-person-orientation when the speaker in fact identifies with another agent *quasi-first-person-orientation*.

A quasi-first-person-orientation is not available with ordinary first-personal pronouns except in certain contexts, such as contexts of imagination, as in (21) (Williams 1973):

(21) I imagine that I am Napoleon.

In (21) the content of the imagination does not involve an identification of the speaker's actual person with Napoleon, but rather the speaker simply projects himself onto Napoleon. Apart from attitudes such as imagination and desire, *de se* interpreted pronouns and other type 2 expressions do not allow for a quasi-first person orientation. The analyses of type 1 and type 2 expressions that I will give later will account for this difference.

⁹ For the examples of *because*-clauses see Egan et al. (2005); for the example with questions see Lasersohn (2005). Of course for those authors the examples support somewhat different accounts.

2.6 Explaining the kind of context dependency involved

A final problem for the standard relativist account is that it does not explain why the phenomena in question give rise to a relativization of truth rather than context-dependency in the traditional sense. The approach makes it look like an accidental fact that evaluative predicates and epistemic modals give rise to an enrichment of the context of assessment, rather than the context of use, unlike, for example, spatial demonstratives like *there*. The account I will propose will explain this difference: relative truth intuitions arise only with expressions that involve an essential first-person orientation, such as evaluative predicates and epistemic modals; they cannot arise with expressions that lack the relevant kind of first-person orientation, such as spatial demonstratives.

The objection just raised does not obtain for all relativist accounts, though, namely not for those that include only the relevant agent in the context of assessment (Egan et al. 2005; Lasersohn 2005).

3 Standard *de se* and type 2 expressions

I will now sketch an account of attitudes *de se* which is close to that of Lewis, but will in addition involve the notion of an attitudinal object, which is needed to account for the truth conditions of sentences with *de se* interpreted pronouns, for conditions on sharing, and for certain differences among type 1 and type 2.

We have seen that like sentences giving rise to intuitions of relative truth, sentences with *de se* interpreted pronouns and other type 2 expressions exhibit sharing. However, there is one difference. While the inference in (22a) is clearly valid, the same inference with (22b) or (22c) as conclusion is valid too:

- (22) a. John thinks that he is the winner.
 Bill thinks he is the winner.
 John and Bill think the same thing.
 b. John and Bill think different things.
 c. John and Bill have different thoughts.

The latter kind of inference is not valid with type 1 expressions:

- (23) John thinks that frog legs taste good.
 Bill thinks that frog legs taste good.
 John and Bill think different things (have different thoughts).

I will call the property of type 2 expressions of allowing both kinds of inferences *optional sharing*.

Let us then look at a simple description of a propositional attitude *de se*:

- (24) John expects PRO to win.

It is generally agreed that for propositional attitudes *de se*, neither a mode of presentation or nor in fact the agent's actual self needs to be part of the propositional

content expressed; requirements that are captured by of Lewis' (1979) account of attitudes *de se* as self-ascriptions of properties, as in the analysis of (24) in (25):

(25) expect(John, $\lambda x[\text{win}(x)]$)

This account predicts sharing if it is the property arguments that are taken to be the objects shared, but it does not account for the optionality of sharing.

Lewis' account must moreover be supplemented by a specification of the truth conditions of propositional contents of attitudes *de se*. Properties as contents of attitudes are truth-conditionally incomplete. But there are clear intuitions that the content of an attitude *de se* has its own truth conditions. If John believes that he is the winner, then what he believes is true or false. What could the truth-conditionally complete object be, the object one makes reference to with what John believes? It is what I call an *attitudinal object* (Moltmann 2003a, b). In the present case the attitudinal object is the object that is John's belief that he is the winner. John's belief that he is the winner is intuitively either true or false. Other attitudinal objects, such as John's expectation to win do not have truth conditions, but corresponding fulfilment conditions. Attitudinal objects arguably are the semantic values of propositional anaphora and of descriptions like *what John believes* (Moltmann 2003a, b). Attitudinal objects should also be considered the objects of the acceptance of assertions. Thus, Joe's assertion that he himself is a hero aims at making the addressee accept that Joe is a hero, not that he, the addressee, is a hero. That is, it aims at making the addressee accept the attitudinal object that is 'the speaker's assertion of PRO being the hero'.¹⁰

In this paper, I make use simply of the intuitive notion of an attitudinal object, referring the reader to other work where I have elaborated that notion and its importance in other respects.¹¹

If attitudinal objects are the semantic values of propositional anaphora, this accounts for the inference with (22b, c) as conclusion. But how could they account for sharing, the conclusion in (22a)? What about the first inference, that is, sharing? There is an entity, closely related to attitudinal objects, that is a suitable semantic value of *the same thing* in the conclusion of (22a), and that is of the sort 'the thought of being the winner'. This entity does not include a particular agent, but can be shared by different agents. It is a *kind of attitudinal object*, a universal whose instances are attitudinal objects of the sort 'John's thought that he is the winner' (cf. Moltmann 2003a, b).

The availability of both attitudinal objects and kinds of attitudinal objects as semantic values of propositional pronouns and quantifiers is what explains optional sharing.

Other type 2 expressions can be treated the same way. While *the museum is to the left* expresses a mere property, propositional pronouns and quantifiers take as

¹⁰ Stalnaker (1981) takes the behaviour of sentences with *de se* interpreted pronouns to be grounds for rejecting Lewis' account of such sentences as expressing properties. Stalnaker instead takes them to express propositions like any other sentences.

¹¹ In Moltmann (2003a), I argue that attitudinal objects should be considered prior to propositions and to be the primary bearers of truth values.

semantic values entities of either the sort ‘John’s belief that the museum is to the left’ or of the sort ‘the belief that the museum is to the left’.

It is now obvious why sentences with type 2 expressions do not give rise to faultless disagreement: the relevant attitudinal objects are truth conditionally complete and they can give rise to disagreement, but not to faultless disagreement.

Type 1 expressions differ from type 2 expressions in that they give rise to faultless disagreement, quasi-*de se* orientation, and obligatory sharing. I would like to argue that these differences are due to:

1. a difference in content: sentences with type 1 expressions have a generic content.
2. a difference in conditions of grasping the propositional content.

These two features are most clearly exhibited by a particular kind of type 1 expression, namely generic *one*, which also displays the intuitions of relative truth. I will show that the semantic analysis of generic *one* which I had proposed in earlier work can and should be extended to other type 1 expressions.

4 First-person-based genericity with generic *one*

Sentences with generic *one* exhibit just the same intuitions of relative truth as sentences with predicates of taste. Generic *one* displays rather transparently what I call *first-person-based genericity*, and first-person-based genericity, I will argue, is the true source of intuitions of relative truth in general.

What is crucial about first-person-based genericity is first that the propositional content of the sentence is generic and second that that content requires a first-person access by whoever maintains or evaluates it. The main difference between the present account and standard relative truth theories is that on the present account it is not the propositional content whose truth is relative to an agent, but rather the cognitive access to the propositional content, which requires an agent to grasp the content in a first-personal way, whatever his evaluative or epistemic background may be.

4.1 Crucial data about generic *one*

Let me first present some basic linguistic data concerning generic *one*.¹² Generic *one* is a pronoun that always leads to generic sentences and it bears a particular connection to the first person: it leads to genericity that is either based on a first-person attribution, as in (26a), or directed towards a first-person attribution, as in (26b):

- (26) a. One can see the picture from the entrance.
 b. One is not allowed to enter.

Generic *one* alternates with arbitrary PRO, which is its empty counterpart and thus occurs in those contexts in which an empty pronominal element, rather than an overt noun phrase is required (Moltmann 2006). Arbitrary PRO and generic *one* can

¹² For a linguistically much more detailed discussion of what follows see Moltmann (2006).

covary in contexts like (27a) (that is, they will take the same semantic values under the relevant assignments), and arbitrary PRO may act as the antecedent for *one*, the possessive *one's* or the reflexive *oneself*, as in (27b):

- (27) a. PRO_{arb} to live a great life is to realize *one's* true potential.
 b. The tailor knows what PRO_{arb} to wear at *one's* own wedding.

Generic *one* always occurs in generic sentences, and as such it can occur in apparently two distinct ways: as genericity-inducing, as in the first occurrence in (28a), and as a bound variable, as in the second occurrence in (28a):

- (28) a. *One* sometimes thinks *one's* life is too short.

In both occurrences generic *one* is best taken to be an expression that introduces a variable subsequently to be bound by a generic quantifier Gn, which is formally associated with a syntactic element in sentence-initial position (Moltmann 2006):

- (28) b. Gn x x sometimes thinks that x's life is too short.

In the linguistic literature on genericity, it is standard to use a generic quantifier Gn as in (28b) (see for example Krifka et al. 1995; Cohen 2002; Kadmon and Landman 1993; Greenberg 2007). There are a lot of controversies, though, of how to understand the generic quantifier (see the overview in Krifka et al. (1995)); in fact how to understand genericity as such (the semantics underlying the generic quantifier) is a major topic not only in linguistic semantics, but also in cognitive science and in philosophy in general. I will restrict myself to adopting just some minimal general assumptions about it. What is generally agreed is that the generic quantifier allows for exceptions and has modal force, that is, it ranges not just over actual individuals, but also over individuals in certain other possible worlds. What individuals the quantifier ranges over is generally driven by conditions of normality or stereotypicality. Such conditions generally are vague and lead to an essentially vague domain (Kadmon and Landman 1993). In addition, it is generally assumed that the generic quantifier is subject to a (much more precise) contextual restriction (Kadmon and Landman 1993). A contextual restriction is needed also for generic *one*. *One* in (29), for example, may range only over the students in a particular class:

- (29) One has to hand in the essay tomorrow.

A plausible, if simplified way of understanding the generic quantifier, quite suited for present purposes, is to take it to be a combination of a universal quantifier ranging over possible worlds, restricted by some accessibility relation R (relating the actual worlds to the 'normal' worlds) and a universal quantifier ranging over individuals which is both restricted by a (vague) condition of normality 'N' and a condition C on contextually relevant individuals:

- (30) $\forall w \forall x (wRw_o \ \& \ x \in D(w) \ \& \ N(w)(x) \ \& \ C(w)(x) \rightarrow P(w)(x))$

Generic sentences thus are heavily context-dependent: the context will provide conditions of accessibility, normality, as well as contextual relevance. At the same time, this context-dependence is subject to certain pragmatic restrictions, such as that the domain of quantification be shared by the interlocutors, and in the particular

case of generic *one*, include the interlocutors themselves. This feature of generic sentences will also play a role for statements of taste to which the analysis will carry over (Sect. 5.1.).

Of course, (30) does not yet capture any first-person-orientation of generic *one*. Before introducing the crucial modification, let us see how the first-person-orientation manifests itself.

The first-person-orientation of generic *one* manifests itself in what, at first sight, appears to be a general availability of an inference to the first person, as in (31):

- (31) One can see the picture from the entrance.
I can see the picture from the entrance.

However, at ‘second sight’, it turns out that this inference is not in fact generally valid. (31) can also be uttered by someone that is for some reason unable to see the picture himself, let’s say someone whose view is temporarily obstructed. Thus, generic *one* may display in fact a *quasi-first-person-orientation*.

The point is made particularly clearly by (32):

- (32) One can see me from the entrance.

(32) does not display any conflict between the grammatical first person and generic *one* because *one* here obviously involves identification with people different from the speaker.

Another particularly clear case of quasi-first-person-orientation is (33), which does not imply that the speaker himself has the mathematical ability to solve the equation:

- (33) One can solve the equation.

The first-person-orientation of generic *one* and arbitrary PRO concerns not only the speaker, but also, in embedded contexts, whoever may be the described agent of the reported attitude or speech act:

- (34) a. John said that one can see the picture from the entrance.
b. John said that it is nice PRO_{arb} to walk in the park.

In attitude contexts, the first-person-orientation of generic *one* is particularly transparent, for example when a generic-*one* sentence is embedded under an epistemic predicate:

- (35) John found out that one can see the picture from the entrance.

For (35) to be true it is sufficient that John has had the experience of seeing the picture from the entrance. Generic-*one* sentences differ thus from universally quantified and other generic sentences, such as (36):

- (36) John found out that people (a normal person) can see the picture from the entrance.

In (36), John has to have made sure in other ways that people other than himself can see the picture from the entrance.

The first-person-orientation of generic *one* manifests itself also in its ability to serve in an immediate description of a first-person experience:

(37) I find that one can easily forget one's own past experiences.

The embedded sentence in (37) naturally serves as a direct description of a first-person psychological state, though the generalizing force is there as well. (37) thus differs markedly from (38), where the attitude described takes as its immediate source third-person observations, or else has a derived content, obtained only inferentially from a first-person experience:

(38) I find that people (a normal person) can easily forget their (his) past experiences.

Generic *one* (and its empty counterpart) is the most suited expression for generalizing subjective experiences as types of experiences.

On the analysis I will give later, the embedded sentences in (37) and (38) do not differ in truth conditions, but rather only in an indication of epistemic grounds. This difference may lead, though, to a difference in truth conditions of an overall epistemic attitude report, such as (37) and (38).

Another manifestation of the first-person-orientation of generic *one* are certain restrictions on which predicates generic *one* can accept, at least on one particular use of generic *one* (I will come to another use below). The predicates *has a nose* and *lives in a big city*, for example, are hardly acceptable with generic *one*, though they are fine in other generic sentences:

- (39) a. ?? One has a nose.
 b. The typical person has a nose.
- (40) a. ?? One lives in a big city.
 b. People live in a big city.

Roughly the restriction on predicates acceptable with generic *one* is that the predicate must describe possible experiences or actions. That is, generic *one* requires predicates whose application to the first person, roughly, requires only self-knowledge; knowledge of one's own experiences, intentions and actions (Moltmann 2006).¹³

First-person-oriented pronouns exhibit both faultless disagreement and sharing. First, faultless disagreement is possible with generic-*one* sentences in just the same way as with predicates of personal taste. One person might be right in asserting (41a), whereas another person, used to a greater level of comfort, may be right in his way in asserting (41b):

- (41) a. One can sleep on this sofa.
 b. One cannot sleep on this sofa.

¹³ There are contexts, however, in which generic *one* imposes no restriction on the predicate whatsoever, namely when occurring as a bound variable, as in (1) and in conditionals as in (2):

- (1) a. Sometimes one forgets that *one* has a nose.
 b. One can doubt that *one* has a soul.
- (2) a. If one lives in a big city, one lives in a city.
 b. If one has a nose, one can breathe.

See Moltmann (2006) for discussion.

Yet the two clearly disagree.¹⁴

Also faulty agreement is possible. Thus, (42) is acceptable even if John and Mary's grounds for their generalization are quite different (lets say when John and Mary have tried out the sofa sleeping in quite different bodily positions or when Mary just found the sofa soft enough and John just long enough to sleep on):

(42) John and Mary agreed that one can sleep on this sofa.

As for sharing, two people, with quite different experiences as their epistemic source, may share the content of a generic-*one* sentence. Thus, an inference of the following sort is always valid, even if, lets say, A's discovery was made by standing at the entrance and B's discovery by seeing a photograph of the entrance:

(43) A discovered that one can see the picture from the entrance.

B discovered that one can see the picture from the entrance.

A and B discovered the same thing (namely that one can see the picture from the entrance).

Free relative clauses and conjunction further support the criterion of 'sharing' with generic-*one* sentences. Thus, (44a, b) are equally possible as conclusions of (43), as is (44c), with a nominalization:

(44) a. A discovered what B found out, namely that one can see the picture from the entrance.

b. A and B discovered that one can see the picture from the entrance.

c. A and B made the same discovery.

Sharing is obligatory with generic-*one* sentences rather than optional. Thus (43) with the conclusion *A and B discovered different things* is not valid.

The parallels between generic-*one* sentences and sentences with predicates of personal taste strongly suggest [1] that the three intuitions of relative truth are to be explained in terms of first-person-based genericity, and [2] that first-person-based genericity is also involved in the semantics of sentences with predicates of personal taste in truth-directed contexts. Let me turn to the latter first.

4.2 The semantic analysis of generic *one*

4.2.1 *The intuitive idea*

The general idea is that sentences with generic *one* as a whole express a generalization based on a first-person application of a predicate, that is, they express first-person-based genericity. First-person-based genericity involves the ability of

¹⁴ Note that the relative-truth intuition about generic-*one* sentences could not be due to the presence of the modal. There is no such intuition about (1a, b) below:

(1) a. Everyone can sleep on this sofa.

b. Everyone cannot sleep on this sofa.

Also, relative truth intuitions arise with generic-*one* sentences without a modal:

(2) a. One often meets celebrities in this café.

b. One does not often meet celebrities in this café.

abstracting from the particularities of one's own person and situation, judging oneself to be normal in relevant respects, and then generalizing to anyone meeting the same conditions. This way of generalizing self-attributions of properties is a form of abstraction, requiring a distinction between relevant and irrelevant features of a given person and his situation. First-person-based genericity can also be viewed as a form of *simulation* in the sense of Gordon (1986, 1995a, b), more precisely as what one may call *generic simulation*.¹⁵ In the case of generic simulation, the relevant intentional agent simply generalizes his own situation, abstracting from the features of his situation that are particular to himself. He does not need to project himself onto a particular other person and make adjustments to adopt the other person's point of view (as in ordinary cases of simulation).

The notion of simulation also helps us understand the quasi-first-person-orientation of generic *one*: first-person-based genericity does not require the agent to actually self-ascribe the predicate. He may just identify with someone to whom he applies the predicate.

Thus, first-person-oriented pronouns involve self-reference that is detached from the relevant agent's actual person: it may involve self-ascribing a property while identifying oneself with someone else and in fact self-ascribing a property while identifying with each one of a collection of individuals.

For the semantic analysis of generic *one* I will make use of a primitive notion of identification 'I', a relation between an agent and another individual with which the agent 'identifies' or whom the agent simulates (or projects himself onto). The basic idea is that generic *one* does not just range over individuals, but individuals as entities the relevant agent identifies with.

4.2.2 The formal analysis

Let us start with the paraphrase of (26a), repeated below as (45a) as in (45b):

- (45) a. One can see the picture from the entrance.
 b. For any one x x as someone with whom the speaker identifies can see the picture from the entrance.

Generic *one* ranges not over individuals as such, but individuals as having a certain property, namely the property of being someone the speaker identifies with (or 'simulates'). There are different ways of construing such entities under a perspective. I will adopt the view that an entity x as having a property P is indeed a different entity than x : it is a qua object in the sense of Fine (1982), namely it is the object x qua being someone the agent identifies with.

On Fine's characterization, qua objects are objects obtained from an individual d and a property P (the 'gloss') such that the following conditions hold:

- (46) For a property P and an individual d ,
1. d qua P exists in a world w at a time t iff P holds of d in w at t .

¹⁵ Gordon's notion of simulation differs somewhat from that of Goldman (1989, 1995), which would not be suitable for the analysis of generic *one*, as discussed in Moltmann (2006).

2. d qua P is identical to a qua object d' qua P' just in case $d = d'$ and $P = P'$.
3. d qua P has a property Q just in case d has Q at the time it is P .

The right notion of a qua object is not quite correctly captured by the third condition given by Fine's characterization: an individual x qua being someone the agent identifies with should have only those properties for whose application the identification provides an epistemic ground or is otherwise relevant, not just any properties that hold of x at the time in question. This corresponds to the actual qualification in ordinary language: John qua being a teacher, or more naturally John as a teacher, cannot have a property like being 35 years old, but as such he may know how children behave, be entitled to a salary, or be competent: properties for which his being a teacher is in some way relevant. The *as*-phrases in (47a) and (47b) provide the epistemic ground for why the predicate holds of the subject:

- (47) a. John as a father knows how children behave.
 b. Jean as a true Frenchman knows about wine.

According to (47a), John's being a father is the reason why John knows how children behave and according to (47b) Jean's being a true Frenchman why Jean knows about wine.

In (47a, b), omitting the *as*-phrase would not make any difference to the truth conditions of the sentence: the *as*-phrase just gives a reason for the holding of the predicate. The proposition expressed by (47a) is quite simply that John knows how children behave and in (47b) that Jean knows about wine. Similarly, in generic-*one* sentences the gloss only serves to provide an epistemic basis for the application of the predicate; it does not affect the truth conditions of the sentence. Moreover, the gloss does not restrict the range of entities generic *one* ranges over. Rather the entities are restricted by both vague conditions on what is considered normal and a contextual restriction, just like the range of any other generic quantifier. The gloss will somewhat influence the domain of quantification, though: the domain will consist of entities the speaker identifies with. As a consequence it is likely to include the speaker as well as the addressee.

With this modification in the definition of a qua object, the restriction to predicates expressing possible experiences or actions follows: the gloss asks for an application of the predicate on a first-person basis even when the predicate is predicated of individuals other than the speaker or relevant agent.

Generic *one* thus introduces a complex variable of the sort ' $\text{qua}(x, \lambda y[I y z])$ '. The crucial modification of the analysis of ordinary generic sentences such as (29b) consists in replacing ' x ' by the two-place functional term ' $\text{qua}(x, \lambda y[I y z])$ ', where the variable ' z ' is to stand for the relevant agent. While, intuitively, the 'ordinary' variable concerns the truth conditions of a generic-*one* sentence, the gloss $\lambda y[I y z]$ gives the 'mode of presentation' that is to govern the applicability of predicates, providing the epistemic basis (or, as we will see, the practical purpose) for applying the predicate. The variable ' z ' will be bound by a lambda operator defining the meaning of a generic-*one* sentence such as *one can see the picture from the entrance* as a property, as in the logical form below:

(48) $\lambda z[\text{Gn } x \text{ can see the picture from the entrance}(\text{qua}(x, \lambda y[\text{I } y \text{ } z]))]$

That is, it expresses the property of being a z such that for any contextually relevant x , x qua being identified with by y can see the picture from the entrance.

Qua-objects are needed not only for the semantics of generic sentences with the pronoun *one*. There are motivations for using qua objects also for other generic sentences. It has been observed in the linguistic literature that singular generic sentences must be ‘definitional’ in nature. Thus (49a) is acceptable because being polyphonic partly defines a madrigal; but not being popular as in (49b) (cf. Krifka et al. 1995; Greenberg 2007):

- (49) a. A madrigal is polyphonic.
b. # A madrigal is popular.

This constraint does not hold for generic sentences with plurals such as (49c), which may express inductive generalizations:

- (49) c. Madrigals are popular.

This indicates that in (49a) or (49b) the predicate applies not to entities x that are madrigals, but rather to qua objects, entities x qua madrigal. *Polyphonic*, expressing an essential property of madrigals, can be true of such a qua object, but not *popular*.¹⁶

Thus, it is a plausible generalization that singular generic sentences, not just those with generic *one*, quantify over qua objects. The gloss of the qua objects in the case of generic *one* comes from the meaning of generic *one*, in the other cases it comes from the restriction of the generic NP. (The effect of the gloss in the case of generic *one* is to guarantee a first-person application of the predicate and thus the basis for the applicability of the predicate will be empirical. As a consequence generic-*one* sentences can express inductive generalizations, whereas in other singular generic sentences, the basis for the application of the predicate can only be inferential.)

How is faultless disagreement possible with generic-*one* sentences? A generic-*one* sentence requires there to be first-personal grounds for asserting or entertaining its content. If one agent has first-personal grounds for asserting the sentence and another for denying it, then both have reasons for their claim (and in this sense neither is at fault), but at most one of them is right in extending the first-personal grounds to anyone in the contextually relevant domain. When speaker and addressee agree or disagree about the truth of a generic-*one* sentence, they disagree about a sentence with agent-independent truth conditions; it is only the epistemic component that generic-*one* sentences carry that will be fixed differently with one agent and with the other. Also the domain of quantification of the generic quantifier should be the same (though as with all generic sentences the domain restriction is essentially vague, and the interlocutors may not have settled what exactly the domain is or indeed how to settle it). The disagreement about generic-*one* sentences

¹⁶ Greenberg (2007) gives a more refined analysis of singular generics in that spirit, without using qua objects. On her account, singular generics involve an ‘in virtue property’, a property that follows from the restriction of the singular generic and that ‘causes’ the predicate to hold.

arises in a completely standard way: it is disagreement about a sentence whose content, if asserted, has the same truth conditions for each interlocutor, truth conditions which are not relative, but absolute.

The intuition of faultlessness of disagreement about a generic-*one* sentence arises because the sentence requires the predicate to be applied in a first personal way, ‘as if to oneself’, to every individual in the domain. Whether the content of a generic-*one* sentence is accepted, rejected, just entertained, or merely understood, this requires the same first-personal access for any agent, that is, an application of the predicate to the individuals in the domain as if to oneself. Formally, the condition that the content of a generic-*one* sentence can be accessed only in a first-personal way consists in a self-attribution of the property expressed by the generic-*one* sentence.

To summarize, a generic-*one* sentence needs to relate to an agent, but not because the truth value of the sentence is relative to an agent, but because the sentence content can be grasped only in a first-personal way. In a case of faultless disagreement, the two agents are not both right in the sense of saying something true. Rather, once the domain of quantification is clarified only one of them or none of them can be right.

Not just the content of a generic-*one* sentence, but also the corresponding attitudinal objects can be grasped only in a first-personal way. If John accepts Mary’s assertion that one can see the picture from the entrance, John needs to self-apply the content ‘that one can see the picture from the entrance’, and if John and Mary share the belief that one can see the picture from the entrance, then that is on the basis of both John and Mary self-applying the content ‘one can see the picture from the entrance’. The role of the property expressed by a generic-*one* sentence in an attitudinal object is different from the role of the property expressed by a sentence with a *de se* interpreted pronoun or another type 2 expression. In the latter case, attributing the property to the agent of the attitudinal object gives the truth conditions of what is believed. In the case of generic *one*, the property must be self-ascribed by whoever grasps the attitudinal object (or its content).

Moreover, in the case of generic *one*, two attitudinal objects with the same property as content are identified even if they involve different agents. Thus, if both John and Mary believe that one can see the picture from the entrance, they can only believe the same thing, not different things, or in other words they can only share the same belief, not have different beliefs (one being John’s belief that one can see the picture from the entrance, the other Mary’s belief that one can see the picture from the entrance). This is why type 1 expressions exhibit obligatory sharing.

4.3 A second use of generic *one*

The gloss of the *qua* objects generic *one* quantifies over not only provides epistemic grounds for the application of the predicate, it may also serve a practical purpose, allowing for an immediate self-application of an independently established generalization (Moltmann 2006). This generalization when using generic *one*, crucially, is presented with the intention of being at least potentially applied in a first-person way by the speaker or, more likely, the addressee, or both. This is the case in particular in deontic sentences:

- (50) a. One is not allowed to enter the room.
 b. The tailor knows what PRO_{arb} to wear.

Whereas with the first-person-based use of generic *one*, the speakers' own experience leads to the generalization expressed by the generic-*one* sentence, in these cases the speaker presents an internalized, but already established generalization, a law, general requirement, or general recommendation. The generalization expressed in (50a), for example, is generally meant to play a role in the addressee's reasoning, so that if the addressee accepts the sentence, this will prevent him from entering the room. Deontic *one* sentences are suited for playing a future role as premises in the addressee's practical reasoning because they involve generalizing self-reference: generic-*one* sentences allow for an immediate first-person application by anyone who accepts them.

As I argued in Moltmann (2006), the two uses of generic *one* do not constitute an ambiguity, but rather correspond to two different strategies of fulfilling one and the same semantic condition imposed by generic *one*, namely that of the gloss generic *one* imposes on the objects quantified over.

5 Generalizing the analysis to other type 1 expressions

First-person genericity, I have tried to show, explains the intuitions of relative truth displayed by generic-*one* sentences. I would now like to show that the same analysis can and should be carried over to sentences with predicates of personal taste. First of all, we will see evidence that sentences with predicates of taste do in fact involve first-person-based genericity. Second, we see that there are also various linguistic links between generic-*one* sentences and sentences with predicates of personal taste, which independently require carrying the analysis of the one over to the other.

5.1 First-person-based genericity and predicates of personal taste

The standard relativist account takes sentences with predicates of personal taste to express a proposition whose truth value depends on the person evaluating the sentence. An assertion of such a sentence ultimately has a purely subjective content: an agent making the assertion should not aim for more than that the sentence be true relative to him.

We have already seen that this makes it hard to explain disagreement in cases of faultless disagreement. I now want to show that on the standard relativist account, the content of a truth-directed act or state involving a sentence with a predicate of taste is as such misdiagnosed: the content is not a 'subjective' one, involving only the agent in question for an evaluation of that content as true; rather it is generic, in fact quantificational (of course with certain restrictions on the domain of quantification).

The first piece of evidence for genericity is the behavior of sentences with predicates of taste in knowledge reports. In general, simple taste statements like *frog legs taste good* and the corresponding first-person knowledge report *I know that frog*

frog legs taste good should be equivalent. First, one direction, the inference in (51), is valid given the rather plausible second premise:

- (51) Frog legs taste good. (uttered by A)
 Whenever something tastes good to me, I know that. (uttered by A)
 I know that frog legs taste good. (uttered by A)

Moreover, the other direction, the inference below, should be valid anyway:

- (52) I know that frog legs taste good. (uttered by A)
 Frog legs taste good. (uttered by A)

The relevant observation now is that the intuition of faultlessness disappears when a simple taste statement is replaced by a knowledge statement:

- (53) A: I know that frog legs taste good.
 B: I know that frog legs do not taste good.

The intuition about (53) is that it is *at most* either A or B that is right, but not both. Both A and B may have grounds for maintaining the content of their knowledge claims, but at least one of them is at fault. On the present account, this follows: the content of A's reported knowledge is in contradiction to B's reported knowledge. Unlike with simple taste statements, the 'fault' of A's and B's assertions in (53) can only concern what is asserted, their knowledge, not what is presupposed, the content of their knowledge. Thus, not the same intuition of faultlessness can arise.

On the standard relativist account, by contrast, both A and B should be right. They are not expected to disagree, though, since their knowledge statements differ in propositional content.

The standard account might appeal to the factivity of *know* to explain the absence of faultlessness in (53). In order for a knowledge statement to be true, the embedded proposition must be true relative to the evaluator (Laserson 2007), and A's and B's statements require that the evaluator have contradictory taste evaluations. The problem with such an explanation, though, is that just the same would hold for non-embedded taste statements which do give rise to faultless disagreement: they also have to be evaluated by a single evaluator.

What distinguishes A's and B's statements in (53) from simple taste statements is, it appears, that the content of the former is not just subjective, but objective. Even if based on a first-person experience, it is a generalization: that frog legs taste good is to be true with respect to anyone of the relevant sort (including the interlocutors). Thus, embedded under a verb of knowledge, a statement of taste expresses as generalization, just like a generic-*one* sentence.

That statements of taste have the same kind of content as generic-*one* sentences can also be seen from the fact that statements like (54a) are (more or less) equivalent to statements like (54b), rather than statements like (54c):

- (54) a. I know that chocolate tastes good.
 b. I know that one likes the taste of chocolate.
 c. I know that I like the taste of chocolate.

The content of a statement of taste embedded under *know* is generic. Though this is not the case with all attitude verbs, namely not for verbs expressing purely subjective attitudes. In English the verb *consider* is particularly interesting in that respect: *consider* requires a predicate of taste in its complement (Lasersohn 2007). But it yields a subjective, not a generic reading.

(55) John considers frog legs tasty.

Also the verb *find*, on one reading, and even *think* can have that interpretation, as in *John find that frog legs taste good* or *John thinks that frog legs taste good* (Stephenson 2007). What is important is that two statements of taste with *consider* (as well as *find* and *think* on the relevant reading) do not give rise to disagreement at all:

(56) A: I consider frog legs tasty.
B: I consider frog legs not tasty.

Here both interlocutors can be right, without disagreeing.

Unlike sentences with predicates of personal taste, generic-*one* sentences do not display any difference in content with verbs expressing truth-directed and with verbs expressing purely subjective attitudes and speech acts:

(57) a. John claims/believes that one can see the picture from the entrance.
b. John thinks/finds that one can see the picture from the entrance.

The content of the embedded sentence in (57b) is as generic as that in (57a).

This supports the view that sentences with predicates of personal taste are ambiguous, depending on the kind of attitude verb under which they are embedded, whereas generic-*one* sentences are unambiguously generic.

The same contrast between truth-directed attitude verbs and verbs expressing purely subjective attitudes holds for speech acts: expressive speech acts (in the sense of Searle 1972) do not give rise to faultless disagreement, only truth-directed ones, such as assertions, do:

(58) a. ? John exclaimed that the wine tastes good. Mary disagreed with him.
b. John claimed that white chocolate tastes good. Mary disagreed with him.

In fact, a simple taste statement like *the wine tastes good* could be either a (truth-directed) assertion or an expressive speech act. Only in the former not in the latter case would disagreement arise.

The difference between the two kinds of propositional attitudes and illocutionary act types has to do with the fact that attitudes like ‘consider’ as well as expressive speech acts do not aim at truth (an aim that can be achieved or fail to be achieved), but consist in the presentation of inherently subjective contents, a content that cannot fail to be true (at least as long as the speech act is sincere). Attitudes and speech acts that aim at truth are incompatible with such purely subjective contents. The two kinds of attitude or speech act verbs, even if they select the same clausal complements, take in fact different kinds of contents as arguments. That is, the ‘purely subjective contents’ of attitudes like ‘consider’ or ‘find’ are of a different sort from the contents that are required by truth-directed speech act and attitude

verbs. The latter are in fact first-person-based generic propositions. *Consider* (in the relevant sense) semantically selects only complements with an evaluative predicate. Evaluative predicates can be construed as two-place relational predicates with one argument position (the judge argument position) being ‘de se’, to be self-ascribed in the attitude of evaluation (Section 5.2.). *Know*, like almost all other attitude verbs, does not select predicates of evaluation with one argument position to be self-ascribed. Instead it generates a generic reading of clausal complements with a predicate of evaluation, which means the judge argument position is filled by a variable bound by the generic operator.

Know of course also allows for evaluative contents in which the ‘judge’ is made explicit:

(59) I believe/claim that frog legs taste good to me.

Frog legs taste good to me is not a subjective, but an objective content, an objective content about a subjective evaluation.

A different sort of evidence against the standard relativist account of predicates of taste and in favour of one in terms of first-person-based genericity comes from factivity in general. Consider a sentence with a predicate of taste embedded under a factive verb such as (60):

(60) John realized that frog legs taste good.

The standard relativist account takes factivity to consist in the requirement that the embedded sentence be true not only relative to the described agent, but also relative to the assessor, such as the speaker himself (Lasersohn 2007). However, there are cases where factivity does not seem to require the statement to be true relative to the speaker. Suppose the speaker does not have any taste sensations, (61) would still have to be true relative to that speaker, which is not possible:

(61) Bill has found out that frog legs taste good.

Factivity should in fact be the condition that the embedded sentence be true for anyone falling under the relevant contextual restriction (relevant and normal individuals), which may exclude the speaker or whoever the assessor may be.

Yet further evidence that sentences with predicates of taste involve the same sort of genericity as generic-*one* sentences is that sentences of the one sort can be used for inferences to sentences of the other sort:

(62) Chocolate tastes good.
One should eat what tastes good.
 One should eat chocolate.

While *tastes good* in the second premise may have the judge variable bound by the generic quantifier explicitly introduced by *one*, this is not so for the first premise. The validity of the inference indicates that the first premise is itself generic, involving the same generic quantifier as generic *one*. Note also that with an explicit judge, as in (63a), or a statement of subjective evaluation, as in (63b), the same inference would not go through:

- (63) a. Chocolate tastes good to me.
 b. I consider chocolate tasty.

The generalization thus is that in truth-directed contexts, sentences with predicates of taste have a first-person-based generic interpretation, the kind of interpretation that is obligatory for generic-*one* sentences in *any* context. Truth-directed propositional attitudes and speech acts go along with a generic interpretation of predicates of personal taste, whereas propositional attitudes and speech acts of individual judgment will go along with a subjective interpretation.

Note that as with generic-*one* sentences, a single first-person experience suffices to make the generalization expressed by a sentence with a predicate of personal taste in a truth-directed context:

- (64) John just found out that chocolate tastes good.

The content of the epistemic attitude described in (64) is a generalization even though what grounds it is most likely a single first-person experience,

Another important similarity between generic-*one* sentences and sentences with predicates of taste is that both allow for quasi-first-person-based genericity.

The problem with the relativist approach in general is that it is unable to distinguish the content of a truth-directed propositional attitude or speech act from a merely subjective content (such as that of the attitude expressed by *consider* or *find*). If a speaker utters *chocolate tastes good* then, knowing the truth-relative semantics of the sentence, the speaker should know that the content of his truth-directed attitude or act would be true just relative to his own context. From his point of view, no considerations need to be made that the content of his utterance also target the context of the addressee. Of course, the speaker may know that the addressee will evaluate the utterance at his context. But why should he be bothered about that and why should it lead to possible disagreement? The addressee, in turn, given his knowledge of the relativist semantics of the sentence uttered should know that too. It thus remains a mystery why the situation should give rise to disagreement. The situation appears entirely undistinguishable from the one where the speaker expresses or upholds his own subjective opinion without targeting the addressee's parameters of evaluation in any way, that is, the situation made explicit by attitude reports like *I consider chocolate tasty*.

5.2 Semantic connections between evaluative predicates and generic *one*

Given that predicates of personal taste have either a subjective or a first-person-based generic interpretation and that obviously the latter is based on the former, the question needs to be addressed: how should the subjective interpretation be represented formally? There are two options for the role of the agent that is involved. First, the judge could be represented by a bound variable; second he may constitute a shiftable contextual parameter. While some authors such as Lasersohn (2005) treat the judge as a shiftable contextual parameter, there are two important reasons to treat him as a value of a bound variable. First this would allow for a straightforward account of the possibility of the variable being bound by a quantified NP, as in (65):

(65) Everyone considers chocolate tasty.

The variable could either be part of an implicit *for*-phrase or occupy an argument position of the predicate of personal taste. The latter option has advantages in accounting for the semantic selection of predicates of taste by verbs like *consider*: *Consider* would select certain two-place predicates with one argument position for a ‘judge’. Let us thus adopt this option. The meaning of *chocolate (is) tasty* in the context of (66a) would then be as in (66b), whereas in the context of (67a) it would be as in (67b):

- (66) a. John considers chocolate tasty.
 b. $\text{consider}(\text{John}, \lambda x[\text{tasty}(c, x)])$
- (67) a. I claim that chocolate is tasty.
 b. $\lambda x[\text{Gn } y \text{ tastes good}(c, \text{qua}(y, \lambda z[\text{I } z \text{ } x]))]$

Sentences with predicates of personal taste in truth-directed contexts (and without an overt subject being specified) always express first-person-based genericity, generalizing from a first-personal subjective experience to anyone meeting the contextually given restriction.

The second reason for positing a ‘judge’-variable for predicates of taste is the possibility of a link between generic *one* and predicates of taste. In fact, this possibility is what gives significant plausibility for the analysis in (47) in the first place. These are two relevant examples:

- (68) a. It is nice when one is walking in the park.
 b. It is nice PRO_{arb} to walk in the park.

In (68a, b), the understood judge of *nice* is the same as, or rather covaries with, the referent of generic *one* and arbitrary PRO. This means that both generic *one* or arbitrary PRO and what is understood as the ‘implicit judge’ of *nice* introduce variables that would have to be bound by a single generic quantifier, as in the analysis of (68a) below:

- (69) $\lambda x[\text{Gn } z \text{ nice}(\wedge \text{walk in the park}(\text{qua}(z, \lambda y[\text{I } y \text{ } x]), \text{qua}(z, \lambda y[\text{I } y \text{ } x]))]$

Also in the contexts below the implicit judge of predicates of personal taste is understood as covarying with arbitrary PRO or else generic *one*:

- (70) a. It is sometimes more pleasant when one walks home than when one drives.
 b. One should walk home because it is so pleasant.
 c. When one drinks this with milk, it is delicious.
 d. When one is young, rollerblading is lots of fun.

The possibility of a single operator binding both the ‘judge’ variable and the variable introduced by generic *one* or arbitrary PRO means that the genericity involved in predicates of personal taste cannot be a matter of the lexical meaning of those predicates (or of *one* of their lexical meanings). Instead it must be tied to the presence of a generic quantifier in the semantic structure of the sentence which is able to bind other elements as well.

5.3 Explaining relative truth intuitions

The explanation of relative truth intuitions with predicates of taste is essentially the same as with generic *one*. Sentences with taste predicates in truth-directed contexts have the same kind of generic content as generic-*one* sentences, a content which as such has agent-independent truth conditions. But just as with generic-*one* sentences, such a content can only be grasped in a first-personal way: understanding, maintaining, or evaluating such a content involves a simulating first-person application of the predicate of taste ‘as if to oneself’. Different speakers will naturally have different grounds for applying or not applying a predicate of taste to other individuals ‘as if to themselves’. But whether an individual *x* is considered as if it was *y* or as if it was *z* will be irrelevant for a predicate of taste *C* to be in fact true of an object *o* and *x*. As in the case of generic-*one* sentences, disagreement arises because of conflicting general claims being made (each with agent-independent truth conditions), and the impression of faultlessness arises because of the first-personal mode of presenting those claims.

This account based on first-person-based genericity also avoids the other problems that at least some standard relativist theories face. First of all, it avoids the meaning-intention problem: no additional parameters such as standards of taste are postulated. The account simply makes use of a first-person attribution of evaluative properties. It moreover makes the right prediction as to which agent serves as the ‘judge’. In independent contexts, the analysis predicts that the judge is the speaker or whoever the speaker may identify with. In attitude contexts, the analysis predicts that the judge is the described agent or whoever the described agent may identify with. Finally, in independent and modal contexts, the judge will be the speaker or whomever the speaker may identify with.¹⁷

The *de se* problem is accounted for because a first-person-based generic sentence expresses a property that is self-ascribed in the speech act or propositional attitude in question. Thus, the judge involved in predicates of evaluation is self-identified by the relevant agent self-ascribing the property expressed by the sentence in question. This is what gives the judge (that agent) a *de se* (rather than a *de re*) status.

To summarize, to accept, entertain, understand, or evaluate the content of a sentence like *Frog legs taste good* can only mean self-applying the property in question. But this self-ascription concerns only the epistemic basis of the claim being made, not its truth. The truth conditions of the claim are agent-independent, though the content of the claim can be grasped only in a first-personal way.

Evaluative predicates raise a particular question concerning their first-person-orientation. There are two distinct ways of applying a predicate to an individual *x*: applying it by simulating *x* (thus a quasi-first-person application) or applying it on the grounds of a ‘theory’ about *x* (a third-person application). The latter is available for the application of evaluative predicates when taking a prepositional phrase that specifies the judge (*tastes good to John*). Without a prepositional phrase, an evaluative predicate obviously must apply either in a *de se* way or, in generic sentences, in a quasi-*de se* way, by applying the predicate ‘as if to oneself’. The second method of

¹⁷ See Koelbel (2002) for discussion and Richard (2004) in particular for relative adjectives such as *rich*.

applying the predicate, a third-person application, is, it appears, unavailable with a simple evaluative predicate. Simple evaluative predicates obligatorily require a self-application or a simulating self-application, in regard to the judge component. For simple evaluative predicates a first-person application is the only option, whereas both options are available for evaluative predicates with prepositional phrase.¹⁸

6 Other type 1 expressions

Until now I have focused on intuitions of relative truth with first-person-oriented pronouns and predicates of personal taste, making use of the notion of first-person-oriented genericity. In what way and to what extent can the account be carried over to other expressions or constructions that have been argued involve a notion of relative truth?

First of all, other evaluative predicates generally exhibit the intuitions of relative truth, that is, moral predicates, aesthetic predicates, and relative adjectives such as *tall* or *rich*.¹³ Though it is less obvious in which sense those predicates express genericity based on a first-person attribution. Any ‘subjective’ application of an aesthetic or moral predicate obviously involves considerations of universal applicability. Thus, such a predicate may involve first-person-based genericity as part of its lexical meaning (unlike predicates of personal taste). At the same time, especially with aesthetic predicates, first-person aspects of meaning (such as triggering emotions) are obviously essential as well. With moral predicates, both emotional aspects of meaning and practical orientation (guidance of actions) require an essential first-person ascription. Perhaps the meaning of predicates of aesthetic or moral evaluation involves first-person attributions of possibly different sorts as part of their more complex meaning. This might also be said for relative adjectives in an evaluative use as discussed by Richard (2004).

In such a (still to be clarified) sense, first-person-oriented genericity would be involved in the semantics of evaluative predicates in general. Note that moral predicates also display the interaction with first-person-oriented pronouns:

- (71) a. It is good PRO_{arb} to treat others with respect.
 b. PRO_{arb} to take the exam is obligatory.

Recall that generic *one* had two uses: one on which the first-person-connection relates to the epistemic grounds for applying the predicate, the other on which the first-person connection gives the practical purpose for applying the predicate. These two kinds of first-person connections, the epistemic and the practical one, are also at play with evaluative predicates. With predicates of moral evaluation the practical connection will clearly be involved, whereas with predicates of taste, the epistemic

¹⁸ An account in the same direction is suggested by Koelbel (2003), who proposes that *chocolate tastes good* be relativised to a ‘perspective’ quite simply because the conditions on the possession of the concept *tastes good* involve a first-person application, which means the content of *chocolate tastes good* could not be grasped otherwise than by a self-application. However, again, unless a perspective is something other than a centred world, it is not obvious how the content of *chocolate tastes good* could be conceived differently on a standard relativist account.

connection will generally be the relevant one. At the same time these generalizations are not strict: predicates of moral evaluation involve not only hypothetical practical reasoning, but also emotions of various sorts (I am outraged to have done X/that you have done X). Similarly, predicates of emotional evaluation will not just involve an inference from a first-person emotional state, but also are to govern future practical reasoning. Even for predicates of personal taste, this may be the case, as Blackburn (1998, Chap. 1) emphasizes. Evaluative predicates thus involve two directions: inference from the first person (generalizing first-person mental states or acts) and inference to the first person (anticipating potential practical reasoning).

The notion of first-person-based genericity is applicable naturally to epistemic modals as well. First, we observe that epistemic modals allow the knower to be made explicit, as in (72):

(72) For all I know, John may be in Paris.

Without a specification of the knower, as in (73a), we seem to have an equivalent to a sentence specifying the knower as 'one', such as (73b):

- (73) a. John may be in Paris.
b. For all one knows, John may be in Paris.

Both (73a) and (73b) give rise to faultless disagreement.

The analysis of epistemic modals that suggests itself would be one according to which *John may be in Paris* could be paraphrased as 'It is compatible with my epistemic state, as far as it is generalizable to anyone of relevance, that John is in Paris'.

With epistemic modals, again the question arises of how to represent the dependence on an agent. Here as well we note that the subject (on which the first-personal content will depend) may be quantificational:

(74) Everyone thinks that John may be in Paris.

It is therefore reasonable to assume that epistemic modals involve an implicit argument position which can be occupied by a variable bound by a generic quantifier. Again using qua objects of the sort 'y qua being someone the relevant agent identifies with', the content of an epistemic modal sentence can be represented as in (75), where M represents the epistemic modal in question:

(75) $\lambda x[\text{Gn } y (M(\text{qua}(y, \lambda z[\text{I } z \text{ } x]), p))]$

Faultless disagreement can now be explained as before: it is disagreement about what from the point of view of the two agents is the common epistemic state. Again the crucial notion is first-person-based genericity, which in this case means genericity based on one's own epistemic state insofar as it is generalizable.

7 Retraction

There is another intuition of relative truth that has been discussed in the literature, in particular in relation to epistemic modals and future contingents. This is retraction

(McFarlane 2005b). Retraction is the possibility of retracting a proposition, once truthfully asserted or believed, at some later point in time. With retraction, it appears one and the same propositional content will give different truth values when evaluated at different points in time. Thus, the content of an utterance of *John may be in France*, correctly asserted at some point in time, may later be withdrawn in view of evidence that excludes John's being in France. Retraction obviously also holds for future contingents.

As Stephenson (2007) points out, retraction does not seem to hold for sentences with predicates of taste, nor does it hold for type 2 expressions. Thus, if John a minute ago believed that the tree is to the left, he will not withdraw his belief once he turns around, maintaining now that the tree is to the right and rejecting his former belief. Rather his former belief content is individuated partly by John's former spatial position, and thus the belief remains true. Retraction thus appears to be a specific form of time-relativity arising with modals (epistemic modals, future contingents). Retraction, as expected, is possible also with generic-*one* sentences containing a modal as in (76).

(76) It is not the case anymore that one can see the picture from the entrance.

But it is equally possible when generic *one* in such a sentence is replaced by an ordinary referential NP, as in (77):

(77) It is not the case anymore that John can see the picture from the entrance.

Obviously, first-person-based genericity is not what is involved in the time-relativity of modals.

8 Conclusion

First-person-oriented genericity is a form of generalization that is associated most explicitly in English with generic *one*. The main point of this paper was that the kind of detached generalizing self-reference expressed by generic *one* is more generally associated with a range of expressions giving rise to intuitions of relative truth. The analysis of those expressions based on first-person-based genericity sheds a very different light on relative truth intuitions themselves: it is not the truth value of the propositional content that must be relativized to an agent, but rather it is the propositional content that must be grasped in a first-personal way to be evaluable as true or false. Such a content once grasped that way has agent-independent truth conditions and thus presents an object of agreement or disagreement of an entirely standard sort.

First-person-based genericity can explain why certain kinds of context-dependent expressions display intuitions of relative truth and only those expressions, namely precisely the expressions whose application conditions involve an essential first-person attribution of a property. It is thus no accident that some context-dependent expressions do not give rise to a relative notion of truth and that conversely evaluative predicates and epistemic modals do not allow for a contextualist or implicit-argument account.

Acknowledgments I would like to thank audiences at the University of Barcelona, New York University, Arché (St Andrews), the IHPST (Paris), and in particular Paul Boghossian and Bob Stalnaker for stimulating discussions.

References

- Blackburn, S. (1998). *Ruling passions*. Oxford: Clarendon Press.
- Cohen, A. (2002). Genericity. *Linguistische Berichte*, 10, 59–89.
- Egan, A. (2007). Epistemic modals, relativism, and assertion. *Philosophical Studies*, 133, 1–22. doi:10.1007/s11098-006-9003-x.
- Egan, A., Hawthorne, J., & Weatherson, B. (2005). Epistemic modals in context. In G. Preyer & P. Peter (Eds.), *Contextualism in philosophy* (pp. 131–169). Oxford: Oxford University Press.
- Fine, K. (1982). Acts, events, and things. In W. Leinfellner, et al. (Eds.), *Sprache und Ontologie. Proceedings of the Eighth Wittgenstein Symposium* (pp. 97–105). Hoelder-Pichler-Tempsky, Vienna.
- Goldman, A. I. (1989). Interpretation psychologized. *Mind & Language*, 4, 161–185. doi:10.1111/j.1468-0017.1989.tb00249.x.
- Goldman, A. I. (1995). In defence of the simulation theory. *Mind & Language*, 7, 104–119. doi:10.1111/j.1468-0017.1992.tb00200.x.
- Gordon, R. M. (1986). Folkpsychology as simulation. *Mind & Language*, 1, 158–171. doi:10.1111/j.1468-0017.1986.tb00324.x.
- Gordon, R. M. (1995a). The simulation theory: Objections and misconceptions. *Mind & Language*, 7, 11–34. doi:10.1111/j.1468-0017.1992.tb00195.x.
- Gordon, R. M. (1995b). Simulation without introspection or inference from me to you. In M. Davies & T. Stone (Eds.), *Mental simulation: Evaluations and applications* (pp. 53–67). Oxford: Blackwell.
- Greenberg, Y. (2007). Exceptions to generics where vagueness, context dependence, and modality interact. *Journal of Semantics*, 2, 131–167. doi:10.1093/jos/ffm002.
- Kadmon, N., & Landman, F. (1993). Any. *Linguistics and Philosophy*, 16, 353–422. doi:10.1007/BF0095272.
- Koelbel, M. (2002). *Truth without objectivity*. London: Routledge.
- Koelbel, M. (2003). Faultless disagreement. *Aristotelian Society*, 104, 53–73. doi:10.1111/1467-9264.t01-1-00003.
- Krifka, M., et al. (1995). Genericity. An introduction. In G. Carlson, F. Pelletier, et al. (Eds.), *The generic book* (pp. 1–124). Chicago: Chicago University Press.
- Lasersohn, P. (2005). Context dependence, disagreement, and predicates of personal taste. *Linguistics and Philosophy*, 28(6), 643–686. doi:10.1007/s10988-005-0596-x.
- Lasersohn, P. (2007). Relative truth, speaker commitment, and control of implicit arguments. In *Proceedings of NELS 37*, and to appear in *Synthese*.
- Lewis, D. (1979). Attitudes *De Dicto* and *De Se*. *The Philosophical Review*, 88, 513–543. doi:10.2307/2184843.
- MacFarlane, J. (2003). Future contingents and relative truth. *The Philosophical Quarterly*, 53, 321–336. doi:10.1111/1467-9213.00315.
- MacFarlane, J. (2005a). Making sense of relative truth. *Aristotelian Society*, 105, 321–339. doi:10.1111/j.1467-9264.2005.00178.x.
- MacFarlane, J. (2005b). The assessment sensitivity of knowledge attributions. In T. Szabo Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology*. Oxford: Oxford University Press.
- MacFarlane, J. (2007). Relativism and disagreement. *Philosophical Studies*, 132, 17–31. doi:10.1007/s11098-006-9049-9.
- MacFarlane, J. (2008). Epistemic modals are assessor-sensitive. In B. Weatherson & A. Egan (Eds.), *Epistemic modality*. Oxford: Oxford University Press.
- Moltmann, F. (2003a). Propositional attitudes without propositions. *Synthese*, 135, 70–118.
- Moltmann, F. (2003b). Nominalizing quantifiers. *Journal of Philosophical Logic*, 35(5), 445–481.
- Moltmann, F. (2006). Generic *One*, Arbitrary *PRO*, and the First Person. *Natural Language Semantics*, 14, 257–281.
- Richard, M. (2004). Contextualism and relativism. *Philosophical Studies*, 119, 215–242. doi:10.1023/B:PHIL.0000029358.77417.df.
- Schiffer, S. (1987). *Remnants of meaning*. Cambridge, MA: MIT Press.

- Schiffer, S. (1990). Meaning and value. *The Journal of Philosophy*, 1990, 602–614.
- Searle, J. (1972). *Speech acts*. Cambridge: Cambridge University Press.
- Stalnaker, R. (1981). Indexical belief. *Synthese*, 49, 129–151.
- Stephenson, T. (2007). Judge-dependence: Epistemic modals and predicates of personal taste. *Linguistics and Philosophy*, 30, 487–525. doi:[10.1007/s10988-008-9023-4](https://doi.org/10.1007/s10988-008-9023-4).
- Williams, B. (1973). Imagination and the self. In B. Williams (Ed.), *Problems of the self* (pp. 26–45). Cambridge: Cambridge University Press.